

MandrakeSecurity
Multiple Network Firewall
User Guide



MandrakeSecurity: Multiple Network Firewall; User Guide

Published 2002-06-15

Copyright © 2002 MandrakeSoft SA

by Camille Bégnis, Christian Roy, Fabian Mandelbaum, Joël Pomerleau, and Florin Grad

Table of Contents

Preface	i
1. Legal Notice	i
2. About Mandrake Linux	i
2.1. Contact Mandrake Community	i
2.2. Support Mandrake	ii
2.3. Purchasing Mandrake Products	ii
3. About This Installation And MandrakeSecurity User Guide	ii
4. Authors And Translators	iii
5. Note From The Editor	iii
6. Tools Used in The Making of This Manual	iv
7. Conventions Used in This Book	iv
7.1. Typing Conventions	iv
7.2. General Conventions	v
1. Getting Started	1
1.1. Getting Started Guidelines	1
1.2. Hardware Requirements	2
2. Installation with DrakX	3
2.1. Introduction to the MandrakeSecurity Installer	3
2.2. Choosing Your Language	4
2.3. License Terms of the Distribution	6
2.4. Disk Detection and Configuration	7
2.5. Configuring your Mouse	8
2.6. Configuring the Keyboard	9
2.7. Selecting the Mount Points	10
2.8. Choose Partitions to Be Formatted	11
2.9. Actual Packages Installation	12
2.10. Root Password	12
2.11. Administrator Password	13
2.12. Adding a User	13
2.13. Configure your Local Network	14
2.14. Where Should You Place the Bootloader	15
2.15. Boot Disk	16
2.16. Installing Updates from the Internet	17
2.17. It's Finished!	18
2.18. How to Uninstall Linux	18
I. MandrakeSecurity Setup And Management	21
3. Basic MandrakeSecurity Setup	21
3.1. Introduction	21
3.2. Basic System Configuration	23
3.3. Configuration of Ethernet Cards	24
3.4. Changing The Administrator's Password	27
3.5. System Log on Local/Remote Machines	27
3.6. Time Configuration	28
4. Configuring Internet Access	31
4.1. Internet Access Status	31
4.2. Analog Modem Configuration	33
4.3. Configure Your ISDN Internet Access	36
4.4. ADSL Connection Setup	41
4.5. Cable/LAN Connection Setup	45
4.6. Provider Accounts Configuration	49
4.7. Time Restriction	50
5. Services: DHCP, Proxy, DNS, And More	51
5.1. Hosted Services State	51
5.2. DHCP Server	51
5.3. Squid Proxy Server	53
5.4. Caching DNS	63
5.5. Intrusion Detection System	64

5.6. Services Activation	64
6. Configuring The Actual Firewall Behavior	67
6.1. Firewall Main Control	67
6.2. Zones Definition	68
6.3. Masquerading, Static NAT and ProxyARP Configuration	74
6.4. Default Policies Configuration	78
6.5. Firewall Rules Configuration	80
6.6. Maintaining the BlackList	84
6.7. Type Of Service Rules Configuration	85
7. VPN Configuration	89
7.1. What is a VPN?	89
7.2. Why a VPN?	89
7.3. Setting up a VPN Server	90
7.4. Setting up a VPN Client	100
8. Configuring Masqueraded Clients	101
8.1. Linux Box	101
8.2. Windows XP Box	103
8.3. Windows 95 or Windows 98 Box	104
8.4. Windows NT or Windows 2000 Box	106
8.5. DOS Box Using the NCSA Telnet Package	110
8.6. Windows for Workgroup 3.11	111
8.7. MacOS Box	111
8.8. OS/2 Warp Box	114
9. Monitoring the Firewall	115
9.1. System and Network Usage	115
9.2. Logs	118
10. Management Tools	127
10.1. Remote Connection Using SSH	128
10.2. Backup and Restore	129
10.3. Update Software	132
II. Applied Theory	135
11. Security Under GNU/Linux	135
11.1. Preamble	135
11.2. Overview	135
11.3. Physical Security	139
11.4. Local Security	143
11.5. Files and File-System Security	144
11.6. Password Security and Encryption	149
11.7. Kernel Security	154
11.8. Network Security	157
11.9. Security Preparation (Before You Go On-Line)	164
11.10. What to Do During and After a Breaking	165
11.11. Security Sources	167
11.12. Frequently Asked Questions	169
11.13. Conclusion	170
Security-Related Terms	171
12. Networking Overview	173
12.1. Copyright	173
12.2. How to Use this Chapter	173
12.3. General Information about Linux Networking	174
12.4. Generic Network Configuration Information	175
12.5. Ethernet Information	179
12.6. IP-Related Information	182
12.7. Using Common PC Hardware	183
12.8. Other Network Technologies	184
12.9. Cables and Cabling	193
A. Where to Find Additional Documentation	197
B. The GNU General Public License	199

B.1. Preamble	199
B.2. Terms and conditions for copying, distribution and modification.....	199
C. GNU Free Documentation License	203
C.1. GNU Free Documentation License	203
0. PREAMBLE	203
1. APPLICABILITY AND DEFINITIONS.....	203
2. VERBATIM COPYING	204
3. COPYING IN QUANTITY.....	204
4. MODIFICATIONS.....	204
5. COMBINING DOCUMENTS	205
6. COLLECTIONS OF DOCUMENTS	206
7. AGGREGATION WITH INDEPENDENT WORKS.....	206
8. TRANSLATION.....	206
9. TERMINATION.....	206
10. FUTURE REVISIONS OF THIS LICENSE	206
C.2. How to use this License for your documents	207

List of Tables

1. Imported Material.....	i
1-1. Hardware Requirements.....	2
12-1. Reserved Private Network Allocations.....	177

List of Figures

2-1. Very First Installation Welcome Screen.....	3
2-2. Choosing the Default Language.....	4
3-1. The Login Window to Connect to MandrakeSecurity.....	21
3-2. MandrakeSecurity Welcome Screen.....	22
3-3. The Log-Out Menu Entry.....	23
7-1. VPN Connection Scheme.....	89
7-2. Adding a VPN Zone.....	91
7-3. Adding an ipsec Network Interface.....	92
7-4. Adding Default Policies for the VPN.....	93
7-5. Adding a Firewall Tunnel.....	94
7-6. Configuring the CA.....	95
7-7. VPN Layout.....	96
7-8. Adding a VPN Server Side.....	97
7-9. Rule to Allow HTTP Traffic Over the VPN.....	98
7-10. Adding a VPN Client Side.....	100
8-1. Reconfiguring the Local Network with drakconnect.....	102
8-2. Setting up the Gateway with drakconnect.....	102
8-3. Setting up The Gateway with Windows XP.....	103
8-4. The Network Icon Under Windows 95.....	104
8-5. The Network Configuration Panel under Windows 95.....	104
8-6. The TCP/IP Configuration Panel under Windows 95.....	105
8-7. The Gateway Configuration Panel under Windows 95.....	106
8-8. The Protocol Configuration Panel under Windows NT.....	106
8-9. The Network Software Panel under Windows NT.....	107
8-10. The TCP/IP Configuration Panel under Windows NT.....	108
8-11. The DNS Configuration Panel under Windows NT.....	109
8-12. Accessing The TCP/IP Control Panel.....	111
8-13. Automatic Configuration of Internet Access For MacOS.....	111
8-14. Manual Configuration of Internet Access For MacOS.....	112
9-1. Sample Snort Report.....	123
10-1. Tools' Main Screen.....	127
10-2. Sample Restore Screen.....	129
10-3. Apply Configuration From The Restored File.....	130
12-1. A Dynamic Routing Example.....	178
12-2. The NULL-Modem Cabling.....	193
12-3. 10base2 Ethernet Cabling.....	194

Preface

1. Legal Notice

This manual (except the chapters listed in the table below) is protected under **MandrakeSoft** intellectual property rights. Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.1 or any later version published by the Free Software Foundation; with the invariant sections being *About Mandrake Linux*, page i, with the front-cover texts being listed below, and with no Back-Cover Texts. A copy of the license is included in the section *GNU Free Documentation License*, page 203.

Front-cover texts:

MandrakeSoft March 2002 <http://www.mandrakesoft.com/>
Copyright © 1999,2000,2001,2002 by MandrakeSoft S.A. and MandrakeSoft Inc.



The chapters quoted in the table below are subject to another copyright owner than the whole manual and a different license:

	Original Copyright	License
"Security Under GNU/Linux", page 135	(c) 1998-2000 Kevin Fenzi and Dave Wreski	GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.
"Networking Overview", page 173	(c) 1997 Terry Dawson, 1998 Alessandro Rubini, 1999 & 2000 Joshua D. Drake {POET}/CommandPrompt, Inc. - http://www.linuxports.com/	LDP License (see Licensing info in chapter)

Table 1. Imported Material

"Mandrake", "Mandrake Linux" and "MandrakeSoft" are registered trademarks of **MandrakeSoft S.A.**; Linux is a registered trademark of Linus Torvalds; *UNIX* is a registered trademark of The Open Group in the United States and other countries. All other trademarks and copyrights are the property of their respective owners.

2. About Mandrake Linux

Mandrake Linux is a *GNU/Linux* distribution supported by **MandrakeSoft S.A.** **MandrakeSoft** was born in the Internet in 1998 with the main goal to provide an easy-to-use and friendly *GNU/Linux* system. The two pillars of **MandrakeSoft** are open-source and collaborative work.

2.1. Contact Mandrake Community

Following are various Internet links pointing you to various **Mandrake Linux** related sources. If you wish to know more about the **MandrakeSoft** company, connect to its web site (<http://www.mandrakesoft.com/>). There is also the **Mandrake Linux** distribution (<http://www.mandrakelinux.com/>) web site and all its derivatives.

First of all, **MandrakeSoft** is proud to present its new open help platform. MandrakeExpert (<http://www.mandrakeexpert.com/>) isn't just another web site where people help others with their computer problems in exchange for up-front fees, payable regardless of the quality of the service received. It offers a new experience based on trust and the pleasure of rewarding others for their contributions.

In addition, MandrakeCampus (<http://mandrakecampus.com/>) provides the *GNU/Linux* community with open education and training courses on all open-software-related technologies and issues. It also gives teachers, tutors and learners a place where they can share knowledge.

There is a site for the “mandrakeholic” called Mandrake Forum (<http://www.mandrakeforum.com/>): a primary site for **Mandrake Linux** related tips, tricks, rumors, pre-announcements, semi-official news, and more. This is also the only interactive web site hosted by **MandrakeSoft**, so if you have something to tell us, or something you want to share with other users, search no longer: this is a place to do it!

In the philosophy of open source, **MandrakeSoft** is offering many means of support (<http://www.mandrakelinux.com/en/ffreesup.php3>) for the **Mandrake Linux** distributions. You are invited in particular to participate in the various Mailing lists (<http://www.mandrakelinux.com/en/flists.php3>), where the **Mandrake Linux** community demonstrates its vivacity and keenness.

Finally, do not forget to connect to MandrakeSecure (<http://www.mandrakesecure.net/>). This site gathers all security related material about **Mandrake Linux** distributions. You’ll notably find there security and bug advisories, as well as security and privacy-related articles. A must for any server administrator or user concerned about security.

2.2. Support Mandrake

By popular request, **MandrakeSoft** proposes that its happy customers make a donation (<http://www.mandrakelinux.com/donations/>) to support the forth-coming developments of the **Mandrake Linux** system. Your contribution will help **MandrakeSoft** provide its users with an ever better distribution, ever safer, easier, up-to-date, and with more supported languages.

For the many talented, your skills will be very useful for one of the many tasks required in the making of a **Mandrake Linux** system:

- Packaging: a *GNU/Linux* system is mainly made of programs picked up on the Internet. These programs have to be packaged so that they will hopefully work together.
- Programming: there are many many projects directly supported by **MandrakeSoft**: find the one that most appeals to you, and offer your help to the main developer.
- Internationalization: translation of the web pages, programs and their respective documentation.
- Documentation: last but not least, the book you are currently reading requires a lot of effort to stay up-to-date with the rapid evolution of the system.

Consult the contributors page (<http://www.mandrakesoft.com/labs/>) to learn more about the way you can contribute to the evolution of **Mandrake Linux**.

On August 3rd 2001, after having established itself as one of the world leaders in Open Source and *GNU/Linux* software, **MandrakeSoft** became the first *Linux* company listed on a European stock market. Whether you’re already a **MandrakeSoft** shareholder or wish to become one, our Investor pages (<http://www.mandrakesoft.com/company/investors>) provide the best financial information related to the company.

2.3. Purchasing Mandrake Products

For **Mandrake Linux** fans wishing to benefit from the ease of on-line purchasing, **MandrakeSoft** now sells its products worldwide from its MandrakeStore (<http://www.mandrakestore.com/>) e-commerce web site. You will find not only **Mandrake Linux** software — operating systems and network tools (Single Network Firewall), but also special subscription offers, support, third party software and licenses, training documentation, *GNU/Linux* related books, as well as other goodies related to **MandrakeSoft**.

3. About This Installation And MandrakeSecurity User Guide

This book includes an introductory chapter, which will guide you in the installation and hardware specifics needed to operate *MandrakeSecurity*. It is advised to read first *Getting Started Guidelines*, page 1 that will give you an overview of the *MandrakeSecurity* life cycle and maintenance tasks.

We will then go through the installation process (“*Installation with DrakX*”, page 3). It is rather straight forward, but if this is your first *GNU/Linux* installation, it is recommended to follow this chapter while you install *MandrakeSecurity*.

Then comes the meat! After this introductory chapter come two parts. The first one is called *MandrakeSecurity Setup And Management* and goes through all the steps needed to run *MandrakeSecurity* efficiently. You will learn the basic setups in the “*Basic MandrakeSecurity Setup*”, page 21 chapter and how to configure your server’s Internet connection in the “*Configuring Internet Access*”, page 31. Then, the “*Services: DHCP, Proxy, DNS, And More*”, page 51 chapter will explain how to configure your server as a DNS, DHCP and Proxy, as well as enable you to use IDS devices.

One of the most important chapters in the first part is the “*Configuring The Actual Firewall Behavior*”, page 67 one. It goes through the Firewall Rules section of *MandrakeSecurity*’s interface and will help you define the inbound/outbound traffic on your network. The last chapters of that first part deal with initial setup and later reconfigurations and tuning of the services. You will find information on those subjects in the “*Monitoring the Firewall*”, page 115 and “*Management Tools*”, page 127 chapters.

The second part is more theoretical, hence its title: *Applied Theory*. It is divided into two chapters. The first one, “*Security Under GNU/Linux*”, page 135, is based on a *HOWTO* by Kevin Fenzi and Dave Wreski. Its main goal is to address security issues system administrators will undoubtedly face. It alternates between philosophical and practical topics on how to better secure your system from potential crackers.

The second and last chapter of the second part is called “*Networking Overview*”, page 173. It is based on a *HOWTO* by Joshua D. Drake {POET}. This chapter contains links to other network-specific documentation on, for example, TCP/IP; it goes through the essentials needed to operate a network properly; it explains technology-oriented principles such as IP and Ethernet-related issues, common technologies to most PC’s, as well as particular network technologies such as Appletalk and Frame Relay.

We conclude this manual with two informative appendices. The first one, *Where to Find Additional Documentation*, page 197, points to information sources on the Internet. And the second one is the *GNU Free Documentation License*, page 203, that is the license which covers the contents of this book.

4. Authors And Translators

The following people contributed to the making of this **Mandrake Linux** manuals:

- Camille Bégnis
- Fabian Mandelbaum
- Roberta Michel
- Rodrigo Pedrosa
- Joël Pomerleau
- Christian Roy
- The people that wrote imported material listed at table 1.

These people also participated at various degrees: Amaury Amblard-Ladurantie, Florin Grad, Philippe Libat, Diane Tan, Hélène Durosini, David Baudens, Philippe Raunet, Philippe Hétroy.

5. Note From The Editor

As you may notice while you go from one chapter to another, this book is a composite document from various authors. Even though much care has been taken in insuring the technical and vocabulary consistency, the style of each author is obviously preserved.

Some of the authors write in English even though it is not their native language. Therefore, you may notice strange sentence constructions; do not hesitate to let us know if something is not clear to you.

In the open-source philosophy, contributors are always welcome! You may provide help to this documentation project by many different means. If you have a lot of time, you can write a whole chapter. If you speak a foreign language, you can help with the internationalization of this book. If you have ideas on how to improve the content, let us know - even advice on typos is welcome!

For any information about the **Mandrake Linux** documentation project, please contact the documentation administrator (<mailto:documentation@mandrakesoft.com>).

6. Tools Used in The Making of This Manual

This manual was written in *DocBook*. *Borges* (<http://linux-mandrake.com/en/doc/project/Borges/>) was used to manage the set of files involved. The XML source files were processed by *openjade* and *jadetex* using Norman Walsh's custom stylesheets. Screen-shots were taken using *xwd* or *GIMP* and converted with *convert* (from the *ImageMagick* package). All this software is available on your **Mandrake Linux** distribution, and all parts of it are free software.

7. Conventions Used in This Book

7.1. Typing Conventions

In order to clearly differentiate special words from the text flow, the documentation team uses different renderings. The following table shows an example of each special word or group of words with its actual rendering and what this means.

Formatted Example	Meaning
<i>inode</i>	This formatting is used to stress a technical term.
<code>ls -lta</code>	Indicates commands or arguments to a command. This formatting is applied to commands, options and file names. Also see the section about “ <i>Commands Synopsis</i> , page v”.
<code>ls(1)</code>	Reference to a man page. To get the page in a <i>shell</i> (or command line), simply type <code>man 1 ls</code> .
<pre>\$ ls *.pid imwheel.pid</pre>	The documentation team uses this formatting for text snapshots of what you may see on your screen. It includes computer interactions, program listings, etc.
<code>localhost</code>	This is literal data that does not generally fit in with any of the previously defined categories. For example, a key word taken from a configuration file.
<i>Apache</i>	This is used for application names. The example used is not a command name but, in particular contexts, the application and command name may be the same but formatted in different ways.
<u>F</u> iles	This is used for menu entries or graphical interface labels in general. The underlined letter indicates the keyboard shortcut, if applicable.
<i>SCSI-Bus</i>	It denotes a computer part or a computer itself.
<i>Le petit chaperon rouge</i>	This formatting identifies foreign language words.
Warning!	Of course, this is reserved for special warnings in order to stress the importance of words; read out loud :-)



This icon highlights a note. Generally, it is a remark in the current context, giving additional information.



This icon represents a tip. It can be a general advice on how to perform a specific action, or a nice feature that can make your life easier.



Be very careful when you see this icon. It always means that very important information about a specific subject will be dealt with.

7.2. General Conventions

7.2.1. Commands Synopsis

The example below shows you the symbols you will find when the writer describes the arguments of a command:

```
command <non literal argument> [--option={arg1,arg2,arg3}] [optional arg. ...]
```

These conventions are standard and you may find them at other places such as the `man` pages.

The “<” (lesser than) and “>” (greater than) symbols denote a **mandatory** argument not to be copied verbatim, but to be replaced according to your needs. For example, <filename> refers to the actual name of a file. If this name is `foo.txt`, you should type `foo.txt`, and not `<foo.txt>` or `<filename>`.

The square brackets “[]” denote optional arguments, which you may or may not include in the command.

The ellipsis “...” mean an arbitrary number of items can be included.

The curly brackets “{ }” contain the arguments authorized at this specific place. One of them is to be placed here.

7.2.2. Special Notations

From time to time, you will be directed to press, for example, the keys `Ctrl+R`, which means you need to press and hold the `Ctrl` and tap the `R` key as well. The same applies for the `Alt` and `Shift` keys.

Also about menus, going to menu item `File→Reload user config` (**Ctrl+R**) means: click on the `File` text displayed on the menu (generally horizontal on the top of the window). Then in the pull-down menu, click on the `Reload user config` item. Additionally, you are informed that you can use the key combination `Ctrl+R`, as described above, to achieve the same result.

7.2.3. System Generic Users

Whenever possible, we used two generic users in our examples:

Queen Pingusa	This user is created at installation time.
Peter Pingus	This user is created afterwards by the system administrator.

Chapter 1. Getting Started

We will review in this introductory chapter all pre-configuration steps required before actually using *MandrakeSecurity* products. Whether you intend to use *MandrakeSecurity* through an appliance or directly from a **Mandrake Linux** installation, this chapter is for you.

1.1. Getting Started Guidelines

The chronological list presented here should guide you through the whole life-cycle of your firewall. Read it carefully before doing anything else, and refer to the cited manual sections as directed.

1. **Hardware Requirements.** If you are building your firewall out of a standard PC, check the adequacy of your hardware for your needs at *Hardware Requirements*, page 1.
2. **Installation.** Install a minimal distribution on the target machine, following the instructions at *“Installation with DrakX”*, page 3.
3. **First connection and basic settings.** Configure the basic system parameters and Internet access: *“Basic MandrakeSecurity Setup”*, page 21.
4. **Services Activation.** Which of the many services proposed by *MandrakeSecurity* do you wish to activate: *“Services: DHCP, Proxy, DNS, And More”*, page 51.
5. **Firewall Rules Setting.** Filter the traffic passing through the gateway: *“Configuring The Actual Firewall Behavior”*, page 67.
6. **VPN Configuration.** If you wish to establish a Virtual Private Network with another remote site equipped with *MandrakeSecurity*: *“VPN Configuration”*, page 89
7. **Client systems setup.** It is now time to connect your different server and hosts to the firewall. Configure the servers in the DMZ according to the firewall rules added in *MandrakeSecurity*. For the clients follow instructions at *“Configuring Masqueraded Clients”*, page 101.
8. **Tests.** Simply make sure the different configured services are working properly. Also test the different firewall rules are actually giving the expected result.
9. **Configuration Backup.** Mandatory, needless to insist: *Backup and Restore*, page 128.
10. **System Monitoring.** Your whole system is now on production and complies with its missions. To make sure everything goes on as expected as time passes, make a good habit of regularly checking the system’s life indicators: *“Monitoring the Firewall”*, page 115.
11. **Changing Passwords.** It is of utter importance to regularly change the admin password used to access your firewall system. To do so access the System Setup Account form from the web interface: *Changing The Administrator’s Password*, page 26.
12. **System Update.** To ensure your firewall is always at the summum of security, **MandrakeSoft** regularly publishes updated packages of applications for which security holes or bugs have been discovered and fixed. Make sure to install the updated packages as soon as they become available: *Update Software*, page 131.
13. **System Deep Reset.** In case it is absolutely necessary: Backup the configuration; Deinstall `naat-*` packages from the system; Install package `snf-en`; restore the configuration.

1.2. Hardware Requirements

If you have chosen to install *MandrakeSecurity* on a standard PC, here are some very rough guidelines regarding the hardware necessary for two different needs. We will then quickly review the installation process.

Configuration	Limited local network with no DMZ and little traffic	Local network plus a DMZ hosting several public Internet servers.
Processor	P166	PIII
RAM	64MB	128MB
Hard Drive	2GB	10GB
Network Interfaces	Ethernet (LAN) + Internet	2*Ethernet (LAN+DMZ) + Internet

Table 1-1. Hardware Requirements

Of course those numbers are purely indicative and are highly dependent on the actual use of the network. Depending on the services actually activated on the firewall, the configuration will have to be upgraded. Check regularly the system load: (*"Monitoring the Firewall"*, page 115) so that you can act before your server actually gets saturated.

Chapter 2. Installation with DrakX

2.1. Introduction to the MandrakeSecurity Installer

DrakX is *MandrakeSecurity*'s installation program. Its ease of use has been enhanced with a graphical user interface, allowing you to move forward and backward through the installation and prompting you when required. With *DrakX*, it doesn't matter whether you're a new user to *MandrakeSecurity* or an old pro – *DrakX*'s job is to give you a smooth installation and an easy transition into *MandrakeSecurity*.



Figure 2-1. Very First Installation Welcome Screen

When you begin, the first screen that comes up will present some information and give you installation options. (figure 2-1). Doing nothing will simply begin the installation in normal or “linux” mode. The next few paragraphs will go over some options and parameters that you can pass to the install program if you run into problems.

Pressing **F1** will open a help screen. Here are some useful options to choose from:

- **vga10**: if you tried a default installation and did not see the graphical interface as shown below in *Choosing Your Language*, page 4, you can try to run the installation in low resolution mode. This happens with certain types of graphics cards, so with *MandrakeSecurity* we give you a number of options to work around problems with older hardware. To try the installation in low resolution mode, type **vga10** at the prompt.
- **text**: if your video card is very old and graphical installation does not work at all, you can always choose the text mode installation. Because all video cards can display text, this is the “installation of last resort”. Don't worry though – it's not likely that you'll need to use the text install.
- **expert**: in some rare cases, your PC may appear to freeze or lock-up during the hardware detection phase. If that happens, then adding the word **expert** as a parameter will tell the install program to bypass hardware detection. Because *DrakX* will not scan for hardware, you will need to manually specify hardware param-

ters later in the installation. The `expert` parameter can be added to the previous modes, so you may end up specifying

```
boot: vga10 expert
```

to perform a low resolution graphical install without *DrakX* performing a hardware scan.



Selecting the `expert` mode will ask you for more details about the installation process, letting you perform a more customized installation.

- `kernel options`: Kernel options usually aren't required for most machines. There are a few cases of motherboards incorrectly reporting the amount of memory installed due to bugs in the design or in the BIOS. If you need to manually specify the amount of DRAM installed in your PC, use the `mem= xxxM` parameter. For example, to start the installation in normal mode with a computer having 256 MB of memory, your command line would look like this:

```
boot: linux mem=256M
```

Now that we've gone over what **might** go wrong, let's move on to the actual installation process. When the installer starts, you'll see a nice graphical interface (figure 2-2). On the left will be the various installation steps. Depending on the installation's progress level, some stages may or may not be available. If a particular stage is available it will be highlighted when you move the mouse pointer over it.

The colors of the buttons on the left side of the screen let you quickly see what's going on with the installation:

- red: this installation phase has not yet been carried out
- orange: the installation stage that is currently being processed
- green: this installation stage has already been configured. However, nothing stops you from going back to a stage that has already been completed if you need to reconfigure something.

This guide assumes that you are performing a standard, step-by-step installation, as described below.

2.2. Choosing Your Language

The first step is to choose your preferred language.




Figure 2-2. Choosing the Default Language

Your choice of preferred language will affect the language of the documentation, the installer and the system in general.


Clicking on the Advanced button will allow you to select other languages to be installed on your workstation, thereby installing the language-specific files for system documentation and applications. For example, if you will host users from Spain on your machine, select English as the default language in the tree view and Spanish|Spain in the Advanced section.

Note that you're not limited to choosing a single additional language. Once you have selected additional locales, click the OK button to continue.




Not all languages listed here are supported in the *MandrakeSecurity* Web interface.

2.3. License Terms of the Distribution



Before continuing, you should carefully read the terms of the license. It covers the entire *MandrakeSecurity* distribution, and if you do not agree with all the terms in it you should click on the **Refuse** button. This will immediately terminate the installation. Clicking on the **Accept** button will continue the installation.

2.4. Disk Detection and Configuration




DrakX will first detect any IDE devices present in your computer. It will also scan for one or more PCI SCSI cards on your system. If a SCSI card is found, *DrakX* will automatically install the appropriate driver.

Because hardware detection is not foolproof, *DrakX* will ask you if you have a PCI SCSI installed. Clicking Yes will display a list of SCSI cards to choose from. Click No if you know that you have no SCSI hardware in your machine. If you're not sure, you can check the list of hardware detected in your machine by selecting See hardware info and clicking the OK. Examine the list of hardware and then click on the OK button to return to the SCSI interface question.

If you had to manually specify your PCI SCSI adapter, *DrakX* will ask if you want to configure options for it. You should allow *DrakX* to probe the hardware for the card-specific options which are needed to initialize the adapter. Most of the time, *DrakX* will get through this step without any issues.

If *DrakX* is not able to probe for the options to automatically determine which parameters need to be passed to the hardware, you'll need to manually configure the driver.

2.5. Configuring your Mouse



Usually, *DrakX* has no problems detecting the number of buttons on your mouse. If it does, it assumes you have a two-button mouse and will configure it for third-button emulation. The third-button mouse button of a two-button mouse can be “pressed ” by simultaneously clicking the left and right mouse buttons. *DrakX* will automatically know whether your mouse uses a PS/2, serial or USB interface.


If for some reason you wish to specify a different type of mouse, select it from the provided list.

If you choose a mouse other than the default, a test screen will be displayed. Use the buttons and wheel to verify that the settings are correct and that the mouse is working correctly. If the mouse is not working well, press the space bar or **Return** key to Cancel the test and to go back to the list of choices.



Wheel mice are occasionally not detected automatically, so you will need to select your mouse from a list. Be sure to select the one corresponding to the port that your mouse is attached to. After selecting a mouse and pressing the OK button, a mouse image is displayed on-screen. Scroll the mouse wheel to ensure that it is activated correctly. Once you see the on-screen scroll wheel moving as you scroll your mouse wheel, test the buttons and check that the mouse pointer moves on-screen as you move your mouse.


2.6. Configuring the Keyboard



Depending on the default language you chose in *Choosing Your Language*, page 4, *DrakX* will automatically select a particular type of keyboard configuration. However, you might not have a keyboard that corresponds exactly to your language: for example, if you are an English speaking Swiss person, you may have a Swiss keyboard. Or if you speak English but are located in Québec, you may find yourself in the same situation where your native language and keyboard do not match. In either case, this installation step will allow you to select an appropriate keyboard from a list.

Click on the More button to be presented with the complete list of supported keyboards.

2.7. Selecting the Mount Points



At this point, you need to decide where you want to install the *MandrakeSecurity* operating system on your hard drive. If your hard drive is empty or if an existing operating system is using all the available space you will have to partition the drive. Basically, partitioning a hard drive consists of logically dividing it to create the space needed to install your new *MandrakeSecurity* system.

Because the process of partitioning a hard drive is usually irreversible and can lead to lost data if there is an existing operating system already installed on the drive, partitioning can be intimidating and stressful if you are an inexperienced user. Fortunately, *DrakX* includes a wizard which simplifies this process. Before continuing with this step, read through the rest of this section and above all, take your time.

If your hard drive has already been partitioned, either from a previous installation of *GNU/Linux* or by another partitioning tool, select the appropriate partitions that you want to install your Linux system into.

If partitions haven't been configured, you will need to create them using the wizard. Depending on your hard drive configuration, several options are available:

- Use free space: this option will perform an automatic partitioning of your blank drive(s). If you use this option there will be no further prompts.
- Use existing partition: the wizard has detected one or more existing Linux partitions on your hard drive. If you want to use them, choose this option. You will then be asked to choose the mount points associated with each of the partitions. The legacy mount points are selected by default, and for the most part it's a good idea to keep them.
- Use the free space on the Windows partition: if **Microsoft® Windows®** is installed on your hard drive and takes all the space available on it, you have to create free space for Linux data. To do so, you can delete your **Microsoft Windows** partition and data (see "Erase entire disk" or "Expert mode" solutions) or resize your **Microsoft Windows** FAT partition. Resizing can be performed without the loss of any data, **provided you previously defragment the Windows partition and that it uses the FAT format. Backing up your data is strongly recommended..** Using this option is recommended if you want to use both *MandrakeSecurity* and **Microsoft Windows** on the same computer.

Before choosing this option, please understand that after this procedure, the size of your **Microsoft Windows** partition will be smaller than when you started. You will have less free space under **Microsoft Windows** to

store your data or to install new software.

- Erase entire disk: if you want to delete all data and all partitions present on your hard drive and replace them with your new *MandrakeSecurity* system, choose this option. Be careful, because you will not be able to undo your choice after you confirm.



If you choose this option, **all** data on your disk will be deleted.


- Remove Windows: this will simply erase everything on the drive and begin fresh, partitioning everything from scratch. **All** data on your disk will be lost.



If you choose this option, **all** data on your disk will be lost.

- Custom disk partitioning: choose this option if you want to manually partition your hard drive. Be careful — it is a powerful but dangerous choice and you can very easily lose all your data. That's why this option is really only recommended if you have done something like this before and have some experience. For more instructions on how to use the *DiskDrake* utility, refer to the online documentation for *DiskDrake* (<http://www.linux-mandrake.com/en/doc/82/en/user.html/diskdrake.html>).

2.8. Choose Partitions to Be Formatted



Any partitions that have been newly defined must be formatted for use (formatting means creating a file system).

At this time, you may wish to reformat some already existing partitions to erase any data they contain. If you wish to do that, please select those partitions as well.

Please note that it is not necessary to reformat all pre-existing partitions. You must reformat the partitions containing the operating system (such as `/`, `/usr` or `/var`) but you do not have to reformat partitions containing data that you wish to keep (typically `/home`).

Please be careful when selecting partitions. After formatting, all data on the selected partitions will be deleted and you will not be able to recover it.

Click on OK when you are ready to format partitions.


Click on Cancel if you want to choose another partition for your new *MandrakeSecurity* operating system installation.

Click on Advanced if you wish to select partitions that will be checked for bad blocks on the disk.

2.9. Actual Packages Installation

Then comes the actual system installation. The packages list is predefined and cannot be changed at this time. The time required to complete the installation depends on the speed of your hardware. An estimate of the remaining time-to-go will be displayed on-screen to help gauge if there is sufficient time to enjoy a cup of coffee.

2.10. Root Password



This is the most crucial decision point for the security of your *GNU/Linux* system: you have to enter the root password. Root is the system administrator and is the only one authorized to make updates, add users, change the overall system configuration, and so on. In short, root can do everything! That is why you must choose a password that is difficult to guess – *DrakX* will tell you if the password that you chose is too easy. As you can see, you are not forced to enter a password, but we strongly advise you against. *GNU/Linux* is as prone to operator error as any other operating system. Since root can overcome all limitations and unintentionally erase all data on partitions by carelessly accessing the partitions themselves, it is important that it be difficult to become root.



The *MSEC* security level is set to 4 (“high”) by default. The password should be a mixture of alphanumeric characters and must be at least 8 characters long. Never write down the root password — it makes it too easy to compromise a system.

One caveat — do not make the password too long or complicated because you must be able to remember it!

The password will not be displayed on screen as you type it in. To reduce the chance of a blind typing error you will need to enter the password twice. If you do happen to make the same typing error twice, this “incorrect” password will have to be used the first time you connect.

2.11. Administrator Password

You are then asked to enter the password for the system administrator (login: *admin*). It is differentiated from the root user, for security reasons, and also because it may not be the same person. It is that *admin* account that will be required to access the *MandrakeSecurity* Web interface. The criterion for choosing this password are the same as for the root password.

2.12. Adding a User

The screenshot shows the MandrakeSecurity - Installation window. On the left is a vertical list of installation steps, each with a circular progress indicator. The steps are: Choose your language, Hard drive detection, Configure mouse, Choose your keyboard, Setup filesystems, Format partitions, Install system, Set root password, Add a user (highlighted with an orange circle), Configure networking, Install bootloader, Create a bootdisk, and Install system updates. The main area of the window displays the 'Enter a user' dialog box. This dialog box has a title bar with a penguin icon and the text 'Enter a user'. It contains the following fields and options: 'Real name' with the text 'Queen Pingusa'; 'User name' with the text 'queen'; 'Password' with a masked field '*****'; 'Password (again)' with a masked field '*****'; and a list of special groups with radio buttons: 'wheel' (selected), 'adm' (allow "su"), 'rpm' (access to administrative files), 'xgrp' (access to rpm tools), and 'access to X programs'. At the bottom of the dialog box are three buttons: 'Accept user', 'Done', and 'Advanced'.

All necessary users have already been added, and you shouldn’t need adding more users for normal *MandrakeSecurity* operations. However, if you plan to use the *squid* PAM authentication feature, you can add here the users that will be authorized.

The first field asks you for a real name. Of course, this is not mandatory — you can actually enter whatever you like. *DrakX* will use the first word you typed in and copy it to the User name field, which is the name this user will enter to log onto the system. If you like, you may override the default and change the username. The next step is to enter a password. From a security point of view, a non-privileged (regular) user password is not as crucial as the root password, but that is no reason to neglect it by making it blank or too simple: after all, **your** files could be the ones at risk.


You can then choose to make that user member of one or more special groups that will give him special privileges. Check the button for the privileges you want for that user.

Once you click on **Accept** user, you can add additional users. Select **Done** when you have finished adding users.



Clicking the **Advanced** button allows you to change the default shell for that user (*bash* by default).

2.13. Configure your Local Network



You will now set up your local network connection (LAN). *MandrakeSecurity* will attempt to autodetect network devices and modems. If this detection fails, uncheck the **Use auto detection** box.




Even though many connection types are offered here, do not configure your Internet connection now. You should now limit yourself to configuring the Ethernet LAN access, so that you can later connect to the administration interface and configure other connections easily through it.

We will not detail each configuration option – just make sure that you have all the parameters, such as IP address, default gateway, DNS servers, etc. from your Internet Service Provider or system administrator.

You will be able to configure all your other network interfaces (Internet, DMZ, etc.) later on through *MandrakeSecurity* interface.

2.14. Where Should You Place the Bootloader




You must indicate where you wish to place the information the bootloader requires to boot to *GNU/Linux*.

Unless you know exactly what you are doing, choose First sector of drive (MBR).

You are then presented the different boot entries that will be proposed at system boot. You can modify them here.

2.15. Boot Disk



The *MandrakeSecurity* CD-ROM has a built-in rescue mode. You can access it by booting the CD-ROM, pressing the F1 key at boot and typing *rescue* at the prompt. If your computer cannot boot from the CD-ROM, there are at least two situations where having a boot floppy is critical:

- when installing the bootloader, *DrakX* will rewrite the boot sector (MBR) of your main disk (unless you are using another boot manager), to allow you to start up with either *Windows* or *GNU/Linux* (assuming you have *Windows* on your system). If at some point you need to reinstall *Windows*, the **Microsoft** install process will rewrite the boot sector and remove your ability to start *GNU/Linux*!
- if a problem arises and you cannot start *GNU/Linux* from the hard disk, this floppy will be the only means of starting up *GNU/Linux*. It contains a fair number of system tools for restoring a system that has crashed due to a power failure, an unfortunate typing error, a forgotten root password, or any other reason.

If you say Yes, you will be asked to insert a disk in the drive. The floppy disk must be blank or have non-critical data on it – *DrakX* will format the floppy and will rewrite the whole disk.


2.16. Installing Updates from the Internet



At the time you are installing *MandrakeSecurity*, it is likely that some packages have been updated since the initial release. Bugs may have been fixed, security issues resolved. To allow you to benefit from these updates, you are now able to download them from the Internet. Choose Yes if you have a working Internet connection, or No if you prefer to install updated packages later.

Choosing Yes displays a list of places from which updates can be retrieved. Choose the one nearest you. A package-selection tree will appear: review the selection, and press Install to retrieve and install the selected package(s), or Cancel to abort.

2.17. It's Finished!



There you are. Installation is now complete and your *GNU/Linux* system is ready to use. Carefully write down the URL given in that dialog, it's the address you'll have to use in your Web browser to access the *MandrakeSecurity* Web interface with admin account. Now, just click OK twice to reboot the system.

2.18. How to Uninstall Linux

The uninstallation process consists of two steps:

1. Delete all partitions on your hard drive and replace them by a single FAT partition with *DiskDrake*.
2. Uninstall the bootloader (generally *grub*) from the Master Boot Record (MBR). To do so, boot under *DOS* and run the `fdisk /mbr` command.

If you have another OS, please consult its documentation to determine how to perform the same step.

Goodbye, and thank you for using *MandrakeSecurity*!

Introducing the MandrakeSecurity Interface

The following chapters are dedicated to the utilization of *MandrakeSecurity*'s web administration tool, which allows you to remotely control your firewall from any of your LAN's machines. The first chapter, "*Basic MandrakeSecurity Setup*", page 21, will guide you through the basic setup of your firewall. You will be able to create accounts, detect and add NICs, set up a syslog server, as well as configure your local time and set up a NTP (Network Time Protocol) server.

Next comes the "*Configuring Internet Access*", page 31 chapter, which will help you configure your server's Internet connection. The third chapter, "*Services: DHCP, Proxy, DNS, And More*", page 51, will enable you to configure services such as DHCP, DNS and Proxy settings. You will also be able to activate an IDS (Intrusion Detection System) device such as *Prelude* or *Snort*, as well as block certain domains or URLs you do not wish your users to visit.

The "*Configuring The Actual Firewall Behavior*", page 67 chapter goes through all the screens included in *MandrakeSecurity*'s Firewall Rules section. Through this section of the web interface, you will also be able to allow/deny traffic between zones.

Finally, we focus on system monitoring (essential to guarantee smooth operation of your firewall system) in the "*Monitoring the Firewall*", page 115 chapter and tools to maintain your system in the "*Management Tools*", page 127 chapter.

We hope you enjoy *MandrakeSecurity*!

Chapter 3. Basic MandrakeSecurity Setup

3.1. Introduction

In this chapter, we will briefly present the interface and how to navigate through it. It is basically made of menus leading to configuration wizards.

3.1.1. Connecting

The connection to the firewall server from any client is made through any modern graphical web browser. The communication is entirely encrypted. Hence, nobody can eavesdrop on the information transferred, especially passwords.

To initiate the session, type in the location field of your browser the URL that was given to you in the last screen of the installation procedure. It should be an address resembling this one:

`https://192.168.1.160:8443/`

where 192.168.1.160 is the IP address of the firewall you chose in the LAN.

You will then get some screens about a new certificate, accept it. Finally, *MandrakeSecurity*'s connection screen appears (figure 3-1).



Figure 3-1. The Login Window to Connect to MandrakeSecurity

Fill it with the admin login and password as defined during the installation. Whenever you are asked to identify to connect to the interface, always use the admin login. The password must be changed the first time you connect, refer to *Changing The Administrator's Password*, page 26.

3.1.2. The Interface



Figure 3-2. MandrakeSecurity Welcome Screen

The interface is designed in a traditional way with a two-level menu on the left and a content frame on the right. The latter will contain the different steps of each wizard corresponding to the second level menu entries. Later on, we will call “section” the topic covered by a first-level menu entry, and “subsection” for second-level menu entries.

Each page of the wizard is made of:

- informative text: what is that screen about;
- user entry fields: to fill or select according to your choices;
- buttons: to perform special actions.

You will also come across icons. These are the most important:

	The Help button. Displays a pop-up window containing help about that particular screen, informing you of the meaning of the various elements present in it.
	The Cancel button. Discards all changes made since the beginning of the wizard and brings you back to MandrakeSecurity's home page.
	The Back button. Makes you go back to the wizard's previous step.
	The Next button. Makes you go to the wizard's next step. Note that the choices made in a page are not validated until the Apply button below is pressed.
	The Apply button. When you reach a wizard's summary last screen, this button allows to confirm the choices and apply them to the system. Do not forget to use it when you have finished a wizard otherwise all your changes will be lost!

3.1.3. Logout

It is very important to explicitly log out of the interface when you are done with all your tasks, or whenever you leave your desk for a certain amount of time. Please note that simply closing the browser is generally not enough as the server has no mean to know that you are leaving your screen unattended, and someone else using your computer just after you may be able to take hold of your session where you left it.




Figure 3-3. The Log-Out Menu Entry

Whenever you finish a session, simply click on that icon. Next time you try to reconnect, you will be asked to identify again.

3.2. Basic System Configuration

This section will help you do a basic setup of your server. It also allows the administrator to change his password to access the interface.

3.2.1. General System Configuration



The information displayed here is very general, yet essential. Your system needs to be associated with a name as well as a domain name. The System and Uptime Info fields give you basic information about your system.

A name will be attributed to the system. That name will then be allocated to a local network. At this point, the parameters to enter depend on whether or not you have a permanent access to the Internet with a fixed IP address.

System Name	firewall.company.net
--------------------	----------------------

This field holds your machine's full hostname: the machine name followed by the domain name, such as firewall.company.net.

Domain Name	company.net
--------------------	-------------

This field holds the machine's domain name. If you hold a domain name and have the required DNSs pointing to your IP address, use it here. Otherwise, use your Internet Service Provider's domain name.

System Info	Linux firewall.company.net 2.4.18-8.1mdksecure #1...
--------------------	--

This field displays, in order, 1) the OS type 2) the machine's full name 3) the kernel version 4) the date it was installed 5) and the processor type.


Uptime Info	4:33pm up 1 day, 23:26, 7 users, load average: 0.00, 0.00, 0.00
--------------------	---

This field displays, in order, 1) the local time 2) the uptime in days, minutes and seconds 3) the number of users 4) and the load average for the past 1, 5, and 15 minutes.

Modify	green check mark
--------	------------------

Click on this check box if you want to change the System Name and/or the Domain Name.


3.2.1.1. System Properties



This section will help you change the System Name and the Domain Name.

When you are done, click on the Next button, then on the Apply button. You will then be brought back to the System Properties Group page, which displays the System Name, Domain Name, System Info and Uptime Info data.

3.3. Configuration of Ethernet Cards



This screen lists the Network Interface Cards (NIC) currently configured on your machine. It will let you select a particular card and reconfigure it, or add another card.

	Zone	IP Address	Subnet Mask	On Boot	Protocol	
eth0	lan			yes	dhcp	edit suppress admin
eth1	lan	10.0.0.1	255.255.255.0	yes	static	edit suppress admin


Each line corresponds to a physical NIC in your computer:

- to reconfigure it, click on the text icon on the left of the trash can. You will also be allowed to select if you want to activate it (or not) at boot time from the Ethernet Interface Configuration page;
- to allow the network associated to this interface to connect to the web interface, click on Admin (see "Administration Interface" below);
- to suppress it, click on the trash can.

Then comes the Administration Interface, which indicates the interface through which administration connections are allowed. This means that your firewall will have to be administered from a computer connected to the sub-network which is associated to the aforementioned card. From it, you can take two actions:

-  - Detect Current NICs: clicking on that icon will launch a NIC auto-detection process. Use it if you previously installed a new NIC in your computer. Note: after your click, it may take some time for the next screen to appear while the computer is detecting new cards.
-  - Add a NIC Manually: should the previous action fail, you can manually configure your card by clicking on that icon.

3.3.1. Detection of Ethernet Interfaces




This screen shows the NIC (or NICs) which has just been automatically detected on your machine. If the card you wish to configure does not appear here, go back to the previous page (using the Back button) and click on the Add a NIC manually button.

```
Driver Mac IP Address Subnet Mask On Boot
Eth0 ne2k-pci 00:40:05:E2:55:F6 192.168.1.160 255.255.255.0 yes
```

Each line corresponds to a physical NIC in your computer. Press the Apply button to confirm the hardware settings.

3.3.2. Ethernet Interface Configuration for your Local Network(s)



In this section, you must define the interface card parameters necessary to satisfy the needs of your local network(s). Some of them may have been chosen already during the installation or a previous configuration and/or filled in with standard values. Make the necessary modifications to answer your present needs.

Connected to the Zone Name	lan
----------------------------	-----

You must choose which kind of network this interface will be attached to. Here are your choices:

- LAN, or systems in your local network(s). These systems must be protected from the Internet and from the DMZ and, in some cases, from each other. Choose this zone to define your local network;
- DMZ, which stands for Demilitarized Zone. Choose this zone if your systems must be accessible from the Internet and from the local network;
- WAN - wide area network. It can be either public or private and it ensures interconnection between computer networks outside of your LAN (e.g. the Internet). Select this type of zone to be connected directly to the outside world.

IP Address	192.168.1.1
-------------------	-------------

Fill out this field if you have a static IP address for that interface. This is your server's address: it is essential since the client systems will refer to that one.

Subnet Mask (ex.: 255.0.0.0)	255.255.255.0
-------------------------------------	---------------

In this field, enter the name of the subnet mask related to the network to which this interface is connected.

Now, set the boot protocol to be used when this interface is initialized. For example, if this interface is connected to the WAN zone, this usually depends on the protocol used by your ISP. Select the appropriate value from the pull-down list on the right:

- Static. This is a permanent IP address assigned to your machine;
- DHCP. This is a dynamic IP address assigned to your machine at boot time. Most cable and DSL ISPs use some form of DHCP to assign an IP to your system. Also, workstations should be set this way to simplify network management;
- bootp. Allows a Linux machine to retrieve its networking information from a server through the network.

Then, you can decide whether or not you want this interface to be activated on each boot.

DHCP Client (optional)	dhcpcd
-------------------------------	--------

This field allows you to choose which kind of DHCP client will be used in your network. You may select one of the following:

- dhcpcd - client daemon which gets an IP address and other information from the DHCP server, automatically configures the network interface, and tries to renew the lease time according to RFC2131 or RFC1541 (the latter, however, is considered obsolete);
- pump - client daemon for BOOTP and DHCP. It enables your machine to retrieve configuration information from a server.
- dhclient - with it, you can configure one or more network interfaces using either the DHCP or BOOTP protocol;
- dhcpxd - its main goal is to conform to the DHCP specification defined in RFC2131. It supports one process per session and is also able to manage all-in-one-process sessions. One of its most advanced features resides in scripts which are run when needed, in order to configure everything required for setting up interfaces.

Finally, you may choose to fill in the DHCP Hostname (optional) field with the appropriate value.

3.4. Changing The Administrator's Password

This form will let you modify the admin login password. It is recommended to change it periodically.

Login Name	admin
New Password	*****
New Password (again)	*****

You need to choose a safe password. Try to select one which includes uppercase and lowercase letters, numbers and special characters, like the interrogation mark (?). When done, click on the Change button.

3.5. System Log on Local/Remote Machines

Logs are an essential part of a security-critical system like a firewall. Not only does it give out information in real time on what is happening on the system, but it also retraces its history, e.g. when something goes wrong in the system -- a crash or an intrusion -- it will find out why it happened and, most generally, figure out a solution.

First of all, you have the choice to activate (or not) the logging system on the local machine (the firewall itself). This, of course, will only be relevant if a display is directly attached to the firewalling machine. It will be possible to control:

Syslog Server (ex: 10.1.1.10)	You can choose to enter either the syslog server's name (i.e.: syslog.company.net) or its IP address. If you don't know the latter, use the ifconfig command as root, or /sbin/ifconfig as a normal user.
--------------------------------------	---

Level for network log	Info
------------------------------	------

This parameter controls the amount of information which will be displayed, according to the level you choose:

- Info: outputs every single message on the firewall, from normal operations to critical messages;
- Notice: returns messages which are not system problematic, but are unusual;
- Warning: informs you that something unusual might be occurring and that you should start thinking of taking action;
- Error: outputs error messages which could lead to system malfunctions;
- Critical: returns messages indicating your system is in serious danger;
- Alert: informs you to take action immediately;
- Panic: outputs only critical messages which generally lead to system failure. At this point, your system might be unusable. Unless you know what you are doing, we strongly advise you not to choose this level.

3.6. Time Configuration

Date (mm/dd/yy)	05/08/2002
Time 24-hour (hh:mm:ss)	08:00:48
Modify	
Time Zone	Europe/Paris
NTP Server Address (optional)	timeserv
Modify	

First of all, the wizard will make two suggestions as to the internal time configuration.



Click on Modify icon () relative to what you want to setup:

- Date and Time: if you do not have a NTP server, click on the Modify button under the Time 24-hour (hh:mm:ss) field to manually set the current date and time on the machine.
- Time Zone and NTP Server Address: click on the Modify button under the NTP Server Address (optional) field to indicate the physical location of the server, and eventually set up a time server, which would automatically set the system's date and time.

3.6.1. Time and Date Setup


Date (mm/dd/yy)	05/08/2002
Time 24-hour (hh:mm:ss)	08:03:22
Apply	Change

Simply enter the current date and time in the respective fields:

Date (mm/dd/yyyy)	10/30/2002
Time 24-hour (hh:mm:ss)	17:07:58

Apply your modifications by clicking on the Change button.

3.6.2. Time Zone and NTP Server Configuration



You need to choose the time zone of your geographical location and, optionally, indicate the presence of an NTP server.

Time Zone	America/Montreal
NTP Server Address (optional)	ntp.time-server.net

In the Time Zone list, select the time zone and then the city closest to you.

Eventually, you can enter the name of a NTP (Network Time Protocol) server, which automatically sets up and checks your clock periodically. If your company has its own server, use it. Otherwise, you can utilize a public server. In that case you can pick one of the list at the Public NTP Secondary (stratum 2) Time Servers web site (<http://www.eecis.udel.edu/~mills/ntp/clock2.html>).

Chapter 4. Configuring Internet Access

In this section, you will be able to configure the way(s) your server will access the Internet. It allows to configure you interfaces with the protocols supported by your version of *MandrakeSecurity*. You can also define all your provider accounts.

4.1. Internet Access Status

Internet Access Group

Current Analog Modem

Current Internet Access Type	PPP
Internet Interface	ppp0
Connection Name	
Modem Port	ttyS1
Modem Speed	57600

Current Provider Account

Provider Domain	
Provider Phone number	
Login Name	
Authentication	PAP
Provider DNS 1	
Provider DNS 2	

Status	On Interface	IP Address
ping:93.161s	eth0	24.201.171.99

Remote Test Host: 198.41.0.6 [Update]

Dynamic DNS Service: off

Dynamic DNS Host Name: [Modify]

[?] Help

Logout

This introductory page to the Internet access configuration wizards summarizes the current Internet access configuration and allows the administrator to manually bring the connection up or down. It also allows to test the connection.

The first part of the frame summarizes all Internet access parameters for the current configuration: type, interface, account information, etc.

Then comes the access status: either Up or Down, and additional information about the current connection. This is followed by three buttons:

- Start: Initiate the Internet connection manually with the current configuration as shown above.
- Stop: Force down the Internet connection.
- Test: update the Internet access status displayed above.

To perform this test, the program simply tries to ping an external host. If you wish to test the connection with a specific host, enter its IP in the Remote Test Host field and click on the Update button.

Remote Test Host	198.41.0.6
-------------------------	------------

It is generally a good idea to use your Internet Service Provider's DNS server as the remote test host.

Dynamic DNS Service	on
Dynamic DNS Host Name	your_host_name.dyndns.org

If you use a dynamic DNS service to configure your host name (i.e. you do not have a fixed IP address), click on the Modify button and complete the corresponding fields to set it up.

4.1.1. Dynamic DNS Setup

This wizard will let you configure a dynamic DNS service to setup your host name

Dynamic DNS Service	www.dyndns.org [dyndns]
Dynamic DNS Host Name	your_host_name
Dynamic DNS Account	your_ddns_account
Dynamic DNS Password	your_ddns_password
Dynamic DNS Notify Mail	admin@company.net

In the first field select your Dynamic DNS Service provider from the pull down list. Then, complete the other fields with the corresponding parameters (host name, username/password and e-mail address) your Dynamic DNS Service provider assigned to you.

Once you are satisfied with your settings, click on the Next button and on the Apply button in the next page to make them effective.

4.2. Analog Modem Configuration

System Setup

Internet Access

- Analog Modem
- ISDN Modem
- DSL Connection
- Cable, Lan
- Provider Accounts
- Schedule

Services

Firewall Rules

VPN

Monitoring

Tools

Logout

Analog Modem Configuration

Current Internet Access Type

Internet Interface

Connection Name

Detect an installed modem

Modem Port

Modem Speed

PPP Special Command

Provider Domain

Provider Phone number

Login Name

Password

Password (confirm)

Authentication

Provider DNS 1

Provider DNS 2

Help

Cancel

Next

This form contains all the parameters required to configure a standard analog modem connection to the Internet. Make sure your ISP provided you with all the necessary parameters.

First, there may be some reminders about the current Internet connection configuration.

Connection Name	My Great Internet Connection
-----------------	------------------------------

Fill this field out with any name that fits the configuration so that you can remember which connection it is relevant to.

You can then try to auto-detect the modem connected to your machine by clicking on the Detect icon:



Detected Modem List	ttyS0
---------------------	-------

This list contains all the modems detected on your machine’s ports (the first serial port in this example). Choose the one you wish to use for this connection.

Modem Port	ttyS1 (COM 2)
------------	---------------

If your modem could not be detected, you can always manually select the port to which it is attached in this list.

Modem Speed	57600
-------------	-------

Simply choose the maximum transfer speed of your modem (in bits/second).

PPP Special Command	
----------------------------	--

In case your connection needs to pass special options to the pppd daemon, you may put them here. You should not need to write anything here in most cases.

Provider Domain	free.fr
------------------------	---------

Enter your Internet Service Provider's domain name.

Provider Phone number	0123456587
------------------------------	------------

The dial-in phone number of your Internet Service Provider. Also enter any dial prefixes needed by the phone connection you are using.

Login name	foo
Password	*****
Password (confirm)	*****

Carefully enter the login name and password as provided by your ISP.

Authentication	PAP
-----------------------	-----

The authentication mechanism used by your ISP. Generally PAP.


Provider DNS 1	123.456.789.122
Provider DNS 2	123.456.789.123

Your ISP's Domain Name Servers.

When you are done with all fields, go on to next step. You will be able to review all parameters and then confirm your choices. The connection will be configured immediately.

If you want to delete a dial-up account, click on the Internet Accounts on the left-side menu.

4.2.1. Analog Modem List




This page displays all modems detected on your machine. Make sure modems are correctly connected and powered on before opening this page.

Detected Modem List	ttyS0 (COM1)
---------------------	--------------

In the pop-down menu, simply choose the port into which the requested modem is plugged, and go on to the next step.

4.3. Configure Your ISDN Internet Access


4.3.1. Choose the Type of ISDN Card



This first step of the ISDN configuration wizard provides you with various options:

- Detect an internal card: The selection of this icon will bring up a list of the ISDN cards detected on your machine. If you have an internal card, try this option first. Otherwise you will either need to:
- Select an internal card: A list of supported ISDN cards will be shown if the previous step has failed.
- Configure an external modem: Select this icon if you have an external modem, not an internal ISDN card.


4.3.2. Choose the ISDN Card



You are here presented the list of the ISDN cards detected on your machine.

Simply select the card you wish to use for your Internet connection and go on to the next step. If your card is not listed, click on "Configure an internal card manually".


4.3.3. Choose the ISDN Card Model



Simply select the name of your card model on the list of suggested models, and go on to the next step.

If your card model is not listed, find out which model is compatible with it in the documentation you were given with your card.

4.3.4. Choose The Provider And Protocol For The ISDN Access



You are here presented with an extensive list of the providers existing around the world. If yours is absent, you will need to configure it manually.

You first need to indicate which protocol to use, this depending on your location:

- Europe
- Rest of the world

Then, you need to configure your provider in one of the following way:

- Select your provider: find your provider in the list which is organized by country, city and finally names of providers. If yours is there, great! just select it and go on to the next step.
- Edit your provider's info manually: If your provider's name did not appear in the above listing, select that icon.

4.3.5. ISDN Access Configuration

This form lists all the parameters required to configure an ISDN connection to the Internet. Make sure you have all the necessary parameters or inquire with your ISP.

If your provider is listed, simply fill in the blank fields.

Your ISDN Login	foo
Your ISDN Password	*****
Your ISDN Password (confirm)	*****

Here, carefully enter the login name and password as provided by your ISP.

Your Personal Phone Number	01.40.41.42.43
-----------------------------------	----------------

What is required here is the number of the phone line you use to connect to the Internet through ISDN.

Provider name	My favorite ISDN provider
Provider Phone Number	01.12.56.89.23

A simple string to identify first your provider and then the phone number you have to dial to connect to the ISDN service of that provider.

Provider DNS 1	123.456.789.122
Provider DNS 2	123.456.789.123

The Domain Name Servers of your ISP.

ISDN Card Description	ELSA Quickstep 1000 (PCI)
------------------------------	---------------------------

This indicates the name of the card being configured.

Dialing Mode	Automatic/manual
---------------------	------------------

Select how you will connect to the Internet:

- Automatic: Whenever the server receives an Internet request compatible with outgoing firewall rules, the connection will automatically be effected.
- Manual: This option will necessitate the intervention of the administrator to manually connect and disconnect that connection when required.


ISDN Card IRQ	12
ISDN Card I/O	0x300

If your card could not be detected, you will need to provide that information. Otherwise leave the fields unchanged.

When all fields have been filled out or left blank as needed, go on to the next step. You will be able to review all parameters and then confirm your choices. The connection will be configured immediately.

4.4. ADSL Connection Setup

4.4.1. Configuration of the ADSL Protocol Type




This is where you can select the specific protocol used by your Internet Service provider (ISP).

Choose the appropriate protocol by clicking on the corresponding check box. If in doubt, ask your ISP.

- Point-to-Point Tunneling Protocol (PPTP)
- Point-to-Point Over Ethernet (PPPOE)
- Dynamic Host Configuration Protocol (DHCP)


4.4.2. Configure A DSL (ADSL) Connection



This is the first screen of the wizard that will guide you through the process of configuring a DSL connection to the Internet. First of all, select the network interface card (NIC) to use for this purpose.

In the list of suggestions, click on the name of the interface you want to use for the Cable/LAN connection. If your specific card seems absent, try detecting it by clicking on the button "Detect".


4.4.3. Add Ethernet Interface



This page shows the interfaces detected on your firewall system.

Simply select the name (ETHx) of the proper card.

4.4.4. Ethernet Interface Configuration For Your Internet Access



What are defined here are the parameters of the interface card necessary to map out the parameters of your xDSL access. Most of the parameters will have been chosen or filled out with standard values already: simply verify that they correspond to your needs.

IP Address (ex: 10.0.0.1)	10.0.0.1
----------------------------------	----------

Fill out this field if you have a static IP address for that interface. Be sure it is the one assigned to you.

Subnet Mask (ex: 255.0.0.0)	255.255.255.0
------------------------------------	---------------

Fill out this field with the subnet mask of the network this interface is connected to. Make sure it is the one you have been assigned.

Default Gateway (ex: 10.0.0.138)	10.0.0.250
---	------------

This is the gateway through which your Internet requests will pass. This parameter is crucial for your firewall machine to reach the Internet.

Finally, you can decide whether this interface will be activated on each boot or not.

4.4.5. Internet Account Configuration For Your DSL Access

The screenshot shows a web-based configuration interface for DSL connection. On the left is a sidebar with a tree view containing: System Setup, Internet Access (selected), Services, Firewall Rules, VPN, Monitoring, and Tools. Under 'Internet Access', there are sub-items: Analog Modem, ISDN Modem, DSL Connection (highlighted), Cable, Lan, Provider Accounts, and Schedule. The main content area is titled 'DSL Connection >' and contains a section 'Internet Account Configuration'. Below this section is a note: 'Please fill the form below (you may not need any of these informations for a DHCP connection)'. The form has several fields: Username, Password, Password (confirm), Provider Name (ex.: provider.net), Provider DNS 1 (ex.: 198.41.0.4), Provider DNS 2, Speed Touch Alcatel Usb Modem (with a dropdown menu showing 'ex:UK'), and Your Country (with a dropdown menu showing 'Belgium'). At the bottom of the form are buttons for '? Help', 'X Cancel', 'Back <', and 'Next >'.

To be authenticated as a user by your provider, you need to give out your account information. The necessary parameters should have been provided by your ISP.

Username	foo
Password	*****
Password (confirm)	*****

Carefully enter the login name and password provided by your ISP. Usually, these are case sensitive.

Provider name	My favorite ADSL provider
Provider DNS 1	123.456.789.122
Provider DNS 2	123.456.789.123

A simple string which first identifies your provider and then the Domain Name Servers of your ISP.

Once all fields are filled out, go on to next step. You will have the opportunity to review all parameters before confirming your choices. The connection will be configured immediately.

4.5. Cable/LAN Connection Setup

4.5.1. Configure a Cable or LAN Connection

Cable, Lan >

Lan or Cable Configuration

Internet Interface	eth0
IP Address	
Subnet Mask	
Default Gateway	
Boot Protocol	dhcp
Activate on each Boot	yes
DHCP Client	/sbin/dhpcd
DHCP Hostname	

External System Name	
External Domain Name	
External DNS 1	
External DNS 2	
Configure your Internet LAN device	Configure
Change your Internet LAN device	Change

This screen appears once an Internet connection of this type is configured. It sums up the current configuration.

- Click on the Change button if you wish to use another NIC for accessing the Internet.
- Select the Configure button if you wish to reconfigure the selected NIC.

4.5.2. Configure a Cable or LAN Connection



This screen will guide you through the process of configuring a cable/LAN connection to the Internet. Those two types of connections are basically identical. First, select the Network Interface Card (NIC) to use for this purpose.

In the list of suggestions, select the name of the interface you wish to use for the cable/LAN connection.

4.5.3. Ethernet Interface Configuration for your Cable/LAN Internet Access

Cable, Lan >

System Setup

Internet Access

- Analog Modem
- ISDN Modem
- DSL Connection
- Cable, Lan
- Provider Accounts
- Schedule

Services

Firewall Rules

VPN

Monitoring

Tools

Ethernet Interface for your Internet Access

Internet Interfaceeth0

IP Address

Subnet Mask

Default Gateway

Optional for DHCP protocol

Boot Protocoldhcp

DHCP Clientdhcpcd

Activate on each BootYes

DHCP HostName

External System Name

External Domain Name

External DNS 1

External DNS 2

? Help

X Cancel

Back

Next

This section will help you define the parameters of the interface card necessary to map the needs of your cable/LAN access. Most of the parameters will already have been selected and the fields filled out with standard values. Start by validating those values.

IP Address	10.0.0.1
------------	----------

Fill this field if you have a static IP address for that interface. Make sure it is the one you have been assigned. Conflicting IP addresses may result in on-going intermittent Internet access problems.

Subnet Mask	255.255.255.0
-------------	---------------

Fill this field with the subnet mask corresponding to the network to which this interface is connected. Make sure it is the one you have been assigned.

Default Gateway	10.0.0.250
-----------------	------------

This is the gateway through which your Internet requests will go through. This parameter is crucial to enable your firewall machine to reach the Internet.

Then, you will have to indicate which boot protocol is to be used when this interface is initialized. This depends on the protocol used by your ISP. Select one of the following:

- static, if you have a specific IP address assigned to your server (most cases for server);
- DHCP, if your address is configured dynamically by this protocol;
- bootp, if your address is configured dynamically by this protocol.

47

Finally, you can decide whether or not to automatically activate this interface every time you boot, the DHCP client and, optionally, the hostname.

Then comes the configuration of your host as a member of the Internet.

External System Name	www.company.net
External Domain Name	company.net

Fill in those fields with the external identification of your firewall machine.

External DNS 1	123.456.789.122
External DNS 2	123.456.789.123

Those DNS IPs generally correspond to your ISP's Domain Name Servers.

4.5.4. Cable/LAN -- Applying the Internet Configuration Changes

Cable, Lan >

Applying the Internet Configuration changes

To accept these values and configure your server, click on "Apply", or use the "Previous" button to correct them

Internet Access	LAN
Internet Interface	eth0
IP Address	
Subnet Mask	
Default Gateway	
Activate on each Boot	yes
Boot Protocol	dhcp
DHCP Client	/sbin/dhpcd
DHCP Hostname	

External System Name	
External Domain Name	
External DNS 1	
External DNS 2	

? Help X Cancel Back Apply

This is the last step needed in order to set up a cable/LAN Internet access.

Review all parameters and click on Apply to publish your modifications, or click on Back if you want to make modifications to your settings.

4.6. Provider Accounts Configuration



This screen presents the current Internet access configuration. In the case that you own several provider accounts, it also allows to switch from one account to another within the same access type.

The first part of the screen informs you about the Internet access currently in use: type, interface, provider. Then comes the list of accounts associated to the current Internet connection type.

DNS1	ProviderPhone	DNS2	Password	Login	Auth
free.fr	123.456.78.1	01010101	123.456.789.2 SecreT	bar	PAP suppress
provider.net	123.456.75.1	02313654	123.456.785.2 SoSecreT	foo	PAP suppress

Each account is made of eight fields:

- Provider Domain: Click on it if you wish to activate this account.
- DNS1: The provider's first DNS server.
- ProviderPhone: If applicable, tells the phone number the modem needs to dial to access the provider.
- DNS2: the provider's second DNS server.
- Password: The password associated to the login account.
- Login: The login corresponding to your provider's account.
- Auth: If applicable, the authentication protocol used to connect to the provider.
- Suppress: Click on this link if you wish to definitely suppress this provider account.

4.7. Time Restriction

In case you have a non-permanent connection, this page will let you define your Internet connection schemes. For each of the three time periods defined, you will be given five options for your connection.

- Dial-up Connect Office: Define the connection schemes during office hours (8:00 AM to 6:00 PM).
- Dial-up Connect Outside: Define the connection schemes outside office hours (6:00 PM to 8:00 AM).
- Dial-up Connect Week-End: Define the connection schemes during the week-end (Saturday, Sunday).

For each of these periods, choose one of the following policies:


- No connection: Connection is down during that period.
- Short Connect Times: Connections are made on demand, and the link cut out whenever requests stop. [relevant only for analog and ISDN modem type links]
- Medium Connect Times: Connections are made on demand, and the link is cut out shortly after requests have stopped. (Irrelevant for permanent type links.)
- Long Connect Times: Connections are made on demand, and the link is cut out much longer after requests have stopped. Average connection delays are thus minimized. (Irrelevant for permanent type links.)
- Continuous Connection: The Internet link is maintained during that period.

When you have gone through the three different time periods, click on the Next button. The following step will show you the choices you made. Review them and go on to the next step, and click on the Apply button to confirm your settings.

Chapter 5. Services: DHCP, Proxy, DNS, And More

This section controls the use of other services, mainly DHCP, DNS and proxy services.


5.1. Hosted Services State



This page displays the state of the services hosted on your machine.

You can change these by clicking on the different services, such as DHCP Server, Web Proxy, Caching DNS and Intrusion Detection, located on the left side of your MandrakeSecurity window.

5.2. DHCP Server




In order to enable the dynamic configuration of new machines connected to your LAN (Local Area Network), you will need to configure a DHCP server on your firewall. When those machines are configured to use a DHCP server on boot-up, they will automatically be set up with all the network parameters they need to

integrate the LAN. Then, you only need to configure the clients to use a DHCP server. This feature is available on most modern operating systems.

Simply choose whether you wish to use (or not) a DHCP server by selecting Yes or No and by clicking on the Next button.

5.2.1. DHCP Server Configuration



You will now need to adjust different parameters in order for your firewall machine to act as a DHCP server for your LAN.

Interface that the DHCP Should Listen to	eth0
--	------

This field holds the name of the interface connected to the LAN. Only those computers which share the same subnetwork with that address will get a response from the DHCP server.

Client's Domain Name (ex.: company.com)	company.com
---	-------------

Simply enter your machine's domain name in this field.

WINS Server IP	If you host a Windows domain name server on your LAN, enter its IP in this field. Hence, the DHCP server will tell your network's Windows workstations, when they boot, what the IP of the WINS Server is, instead of having to configure each Windows workstation accordingly.
----------------	---

Start of the IP Range (ex.: 24)	65
End of the IP Range (ex.: 75)	254

Those fields contain the IP address range allowed for client DHCP hosts. The example given is for a class C subnetwork. Make sure you do not include the first IP (0 in that case) nor the last (255) in the range: they are reserved. Note that the first addresses are generally reserved for static IP hosts, while the last ones are used by DHCP servers.

Default Time Interval (21600 = 6h)	21600
------------------------------------	-------

Max Time Interval (43200 = 12h)	43200
---------------------------------	-------

The assignment of an IP to a host is always limited in time. When the client does not set the needed leasing period, the server will intervene and reassign an IP to the host every "Default Time Interval" seconds. However, a client's request for a specific leasing period inferior to the "Max Time Interval" will be honored. Otherwise, an IP will be reassigned automatically after that "Max Time Interval" has expired.

5.2.2. DHCP Server -- Confirming your Changes

Interface That The DHCP Should Listen to	
eth0	
Client's Domain Name	borges.net
WINS Server IP	
Start of the IP range	65
End of the IP range	254
Default time Interval (21600 = 6h)	21600
Max Time Interval (43200 = 12h)	43200

Buttons: ? Help, X Cancel, Back, Apply

At this stage, the Confirmation Page displays the data which will be applied to the firewall.

If you click on the Apply button, all your changes will be saved, replacing the default or previous settings. If you wish to change the parameters, click on the Back button, change the parameters, click on the Next button, and apply your changes.

5.3. Squid Proxy Server

Select the Proxy Server mode: manual with authentication

Current Web Filtering Status

Url Filter	on		
Content Filter	off		

Buttons: ? Help, X Cancel, Next


To be able to cache HTTP and FTP requests made from inside your LAN to the Internet, you will need to set up a proxy server on your firewall. This allows for a page, which is requested by two different users, to be retrieved only once from the Internet, thus dramatically fastening access to this page while saving precious bandwidth.

MandrakeSecurity has chosen the Squid proxy server. The latter acts as an agent, accepting requests from clients (such as browsers) and passing them on to the appropriate Internet server. It then stores a copy of the returned data in an on-disk cache.

Choose between four options before going on to the next step:

- Deactivate the Proxy Server: if you choose not to use the proxy, requests from users will be directly forwarded to the outside;
- Activate Transparent Proxy: activates the proxy and configures it to act as a transparent proxy, i.e. users will not need to configure their clients in order to use the proxy: all requests are automatically intercepted and managed by the proxy;
- Activate Manual Proxy: same as previous, but client web browsers will need to be reconfigured to explicitly use the proxy server installed on your MandrakeSecurity server;
- Activate Manual Proxy with User Level Authentication: same as previous, but users will be asked for a user name and password to be able to use the proxy. WARNING: create accounts on the firewall Linux box for the users who are authorized to connect to the Internet.

5.3.1. Proxy Main Configuration



This where you can set your proxy parameters. After deciding on a few common parameters, you have the option to activate the web filtering or not.

Squid Mode:	manual
--------------------	--------

This field is defined by the Proxy Server Mode you previously chose between Disable, Transparent, Manual or Manual with Authentication. In our example, we chose Manual.

Squid Port (we recommend 3328)	3328
---------------------------------------	------

This is the port on the firewall machine on which Squid will listen for requests. There is no need to make any changes here unless this port is to be used by another service.

Squid Cache Size (in MB)	100
---------------------------------	-----

This field lets you control the amount of cached data Squid can store and manage. In order for your cache to be efficient, you need to adjust the cache space to the number of users: more users, more space needed. It may vary between 10 MB and 10 GB or more.

Select the Authentication Mode	Pam
---------------------------------------	-----

This field will only be displayed if you selected the Manual with Authentication mode. It allows you to choose between PAM (Pluggable Authentication Modules), a flexible mechanism for authenticating users (the default action with MandrakeSecurity); LDAP (Lightweight Directory Access Protocol), which enables access to on-line directory services; Samba, which will connect you to a Samba workgroup such as COMPANY (if you use such a server); and finally NIS (Network Information Service), which facilitates the communication of critical information to every machine throughout a network.

Squid Admin E-mail	root@company.com
---------------------------	------------------

Type the administrator's e-mail in this field (root@company.com in our example), so your users know to whom to refer bugs/problems (if any).

Once you are back on the main Squid proxy server page, you can activate the web filtering. This feature will enable you to deny or restrict access to certain pages on the Internet, depending on their URLs. It can be useful to block access to ad banners or adult contents.

You can filter by URL or content. Select the text icon for either one. Then, you can set the filtering rules for Authorized Networks, Time Restriction, Advertising to Be Removed, Banned Destination URLs, Privileged IPs, Banned Source IPs, or backup/restore your data.

5.3.1.1. LDAP Options

The screenshot shows the 'System Setup' tab in the Squid Admin interface. Under the 'Services' section, 'Web Proxy' is selected. The 'Authentication Mode' is set to 'Ldap'. The following fields are filled in:

- ou (ex.: people): department
- dc (ex.: mdk): company
- dc (ex.: com): us
- Ldap Server IP: 192.168.2.78

At the bottom, there are buttons for '? Help', 'X Cancel', 'Back', and 'Next'.

In order to use LDAP, you need to set basic parameters such as the ones which follow.

ou (ex.: people)	First, set the object attribute to the category needed, such as "people" or "department".
dc (ex.: mdk)	Next, set the dc object variable to your domain name, such as "mandrake". The other dc object variable should be set to your country, such as "us" for USA.
LDAP Server IP	Finally, type in the LDAP server IP (i.e.: 192.168.2.78) or its name (i.e.: ldap.mesd.k12.or.us).

5.3.1.2. Samba Options

The screenshot shows the 'System Setup' tab in the Squid Admin interface. Under the 'Services' section, 'Web Proxy' is selected. The 'Authentication Mode' is set to 'Samba'. The following field is filled in:

- Samba Workgroup (ex.: MANDRAKESOFT): COMPANY

At the bottom, there are buttons for '? Help', 'X Cancel', 'Back', and 'Next'.

This page allows you to set the name of your Samba Workgroup.

Samba Workgroup (ex.: MANDRAKESOFT)	MY_WORKGROUP
--	--------------

Simply enter the name of your Samba Workgroup in the blank field and click on the Next button.

5.3.1.3. NIS Options

The NIS Authentication page allows you to set two essential parameters: the NIS domain and the NIS list.

NIS Domain (ex.: yp.mandrakesoft.com)	In this field, type in your NIS domain name in the following format: yp.mandrakesoft.com, that is "yp" for Yellow Pages, followed by your domain name.
NIS List (usually yp_list.byname)	Same principle for the NIS list: start with "yp", followed by the underscore sign (" _"), then the list name in the "list.name" format.

5.3.2. Web Proxy Filtering URLs


You have activated Proxy Guard, and this page will let you configure it. This is the wizard's first screen. It will make suggestions to configure the various filtering aspects.

First, check the section to be configured and go on to the next step:

- **Authorized Network:** Enter the authorized networks/masks in order for them to be able to use the proxy's services.
- **Time Restriction:** Allows you to define the connection schedule, i.e. when people are allowed to connect or not.
- **Advertising to Be Removed:** Enter the URLs or whole domains of advertising sites. The images proceeding from those sites will not be forwarded to the clients.
- **Banned Destination URLs:** Enter the URLs or whole domains for which all access should be blocked.
- **Privileged IPs:** Enter the IPs of your local network's privileged machines. They will be freed from any restrictions imposed by the filter to other hosts.
- **Banned Source IPs:** Designates those machines which are not authorized to use the proxy.
- **Backup Restore:** Enables you to make a copy of your proxy rules, and to restore them.

Whenever you are finished configuring one of those preceding sections, you will be brought back to this page. Once you are satisfied with your settings, click on the Apply button to make them effective.

5.3.2.1. Authorized Source Network/Mask



This form will let you specify which subnetworks are allowed to use the proxy services. If different classes of machines are to be specified on your network (one for people authorized to access the web, the other for unauthorized people), you will have to create a subnetwork for the authorized machines and assign IPs in accordance with the authorization status the particular computer has been granted.

Authorized Network/Mask Entry	192.168.1.0/25
--------------------------------------	----------------

In this field, enter the Network/Mask IP address of the subnetwork (the example shown designates the IP range from 192.168.1.0 to 192.168.1.127 -- where 0 is your network address and 127 your broadcast address).

Then click on the Add button:  The address will appear in the list at the bottom of the page.

If you wish to suppress an IP from the list, simply select it and click on the Suppress button:



When you have gone through the list, go on to the next step which will bring you back to the main web-proxy filtering page.

5.3.2.2. Time Restriction

Time Restriction

Do you want to: enable it

Define the authorized time zones. Note that each day is split into two zones: AM and PM.
Example: Sun AM 09h-12h and Sun PM 13h-22h

Sun AM	09:30-12:00	Sun PM	13:00-19:00
Mon AM	09:00-12:00	Mon PM	13:00-19:00
Tue AM	09:00-11:00	Tue PM	12:00-19:00
Wed AM	09:00-12:00	Wed PM	12:00-18:00
Thu AM	09:00-13:00	Thu PM	13:00-18:00
Fri AM	09:00-12:00	Fri PM	13:30-18:00
Sat AM	08:20-13:00	Sat PM	13:30-19:00

? Help X Cancel Back Next

This form will let you define time periods within which access to the proxy will be allowed. Note that this does not affect privileged hosts of your local network. Outside these time periods, restricted machines will not be able to browse the web.

You first need to choose to enable or disable this feature. If you enable it, you will have to define the time periods in question: there are two periods per weekday.

Sun AM	09:00-13:00
--------	-------------

Make sure to strictly respect the time format as illustrated: HH:MM-HH:MM. Modify all periods at your convenience.

When you are done with all periods, go on to the next step. You will be shown what choices you have made. Review them and go on to the next step. This will bring you back to the main web-proxy filtering page.

5.3.2.3. Advertising Domains/URLs

New banned advertising Domain:

Add this Domain to the list

New banned URL:

Add this URL to the list

Suppress selected entry

Suppress selected entry

Domains

- ads.freshmeat.net

URLs

- us.a1.yimg.com/us.yimg.com/a/an/anchor

How can we get rid of those bothersome ads on our favorite web sites? Let's examine two examples: freshmeat.net and yahoo.com.

On the freshmeat.net web site, right-click on the advertising picture and then on Copy image location. Then go to the New Banned Advertising Domain section, click on the middle mouse button or on both mouse buttons, simultaneously. Erase the last part of the URL and you will get ads.freshmeat.net. Now, add this domain to the list.

New Banned Advertising Domain	ads.freshmeat.net
--------------------------------------	-------------------

On the yahoo.com web site, right-click on the advertising picture and then on Copy image location. Then go to the New Banned Advertising URL section and click on the middle mouse button or on both mouse buttons, simultaneously, to paste the information. Erase the last part of the URL and you will get us.a1.yimg.com/us.yimg.com/a/pr/promo/anchor. Now add this URL to the list. If you click on your browser's Refresh button several times and copy the image location, you will get several different URLs:

```
us.a1.yimg.com/us.yimg.com/a/ya/promo/anchor
us.a1.yimg.com/us.yimg.com/a/an/promo/anchor
us.a1.yimg.com/us.yimg.com/a/ya/yahoopager/messenger
us.a1.yimg.com/us.yimg.com/a/ya/yahoo_auctions
```

and so on...

New Banned URL	us.a1.yimg.com/us.yimg.com/a/an/anchor
-----------------------	--


The ads corresponding to this particular URL will not be displayed.

The lists for the two categories will then appear. Select any item in those lists and delete it by selecting it and

clicking the Suppress button:  , if you wish to do so.

When you have gone through those lists, go on to the next step. This will get you back to the main web-proxy filtering page.

5.3.2.4. Banned Sites



This form offers you three different ways to filter the pages viewed within the local network. These three types of filtering depend on those pages' URLs.

New Keyword	microsoft
--------------------	-----------

All URLs containing this word will be blocked. Click on Add This Keyword to the List to do so.

New Banned Domain	msn.com
--------------------------	---------

All pages dependent on a server whose name ends with this domain will be blocked. In our example: "http://eshop.msn.com/category.asp?catId=212" will not be displayed. Click on Add this Domain to the List to do so.

New Banned URL	www.XXX.com/index_ns.html
-----------------------	---------------------------

This specific URL will not be displayed. In our example:

http://www.XXX.com/ad.html

won't be discarded. Click on Add this URL to the List to do so.

Then, the lists for the three categories will be shown. You can select an item on those lists and delete it by



selecting it and clicking the Suppress button:

When you have gone through those lists, go on to the next step, this will bring you back to the main web-proxy filtering page.

5.3.2.5. Privileged IPs

This form will let you add or remove the IPs of the privileged machines of your local network. Those machines will be freed from any restrictions imposed by the filter to other hosts.

Enter a New Privileged IP Address	192.168.1.111
--	---------------

Enter the full IP address of the privileged host. Then click on the Add button:



The IP will appear in the list at the bottom of the page.

To suppress an IP's privileges on the list, simply select it and click on the Suppress button:



When you have gone through that list, go on to the next step. This will bring you back to the main web-proxy filtering page.

5.3.2.6. Banned IPs

This form will let you add or remove the IPs of those machines not authorized to use the proxy at all. This means that if, from the local network, there is no other gateway to the Internet, these machine users will not be able to browse the web.

Enter a New Banned IP Address	192.168.1.110
--------------------------------------	---------------

Enter the full IP address of the banned host. Then click on the Add button:




The IP will appear in the list at the bottom of the page.

If you wish to suppress an IP from the list, simply select it and click on the Suppress button:



When you have gone through that list, go on to the next step. This will bring you back to the main web-proxy filtering page.


5.3.2.7. Doing Backups and Restoring Proxy Rules



Doing a backup of your proxy rules is an excellent idea. Being able to restore them is also very convenient. Follow these easy steps to do so.

- **Backup:** to create a backup of your proxy rules, simply click on the grey Backup button. A new page will be displayed: Shift-click on the WebProxy Backup File link (or right-click and choose Save Link As) and the WebProxyRulesBackup.tar.bz2 file will be saved to your hard disk.
- **Restore:** in order to restore your proxy rules, click on the Browse button to locate the appropriate file (such as WebProxyRulesBackup.tar.bz2). Then, click on the Upload button and, on the next page, on the Apply button.

5.3.2.8. Backup Rules




Your backup operation was succesful. Now, follow the instructions below to save your proxy rule backup file.

Simply shift-click on the WebProxy Backup File link (or right-click and choose Save Link As) and the WebProxyRulesBackup.tar.bz2 file will be saved to your hard disk. Then click on the Next button.

However, if you wish to terminate this operation, click on the Cancel button: you will be brought back to the MandrakeSecurity main page.

5.4. Caching DNS




This wizard will help you configuring the DNS server.

DNS is the acronym for Domain Name System. It translates human-readable machine names into machine-readable IP addresses and vice-versa. This configuration wizard will provide a local DNS service for computers connected to your local network and all non-local requests will be forwarded to an outside DNS server.

To activate this service, please check the Enable box and click on the Next button.

Clicking on Cancel button will take you to the default (starting) page cancelling this wizard.

5.4.1. DNS Forwarders Configuration Step



Here you can specify the DNS servers used for forwarding requests to.

DNS Server	on
Primary forwarder DNS	your_ISPs_PrimaryDNS_IP_address
Secondary forwarder DNS	your_ISPs_SecondaryDNS_IP_address

You need to enter the primary and, optionally, secondary DNS servers' IP addresses into the corresponding fields. Usually, you will enter your ISP's primary/secondary DNS servers' IP here, but you might be using another DNS server to forward requests to.

Once you have entered the IP address(es) for your DNS forward server(s), click on the **Next** button.

Clicking on **Back** will take you back to the previous step of this wizard.

Clicking on **Cancel** will take you to the default (starting) page cancelling this wizard.

5.5. Intrusion Detection System

Intrusion Detection System

Activate / Deactivate Intrusion Detection Systems

! Warning ! These options could generate huge logs

Make your choices, and Press 'Next' to confirm your modifications.

Prelude IDS: ☒ Enable

Snort IDS: ☐ Enable

? Help X Cancel ➔ Next

This page allows you to select and activate an Intrusion Detection System (IDS) on your server. IDS are signature based packet filtering applications used to generate alarms upon abnormal network activities.

Prelude IDS	Enable/Disable
Snort IDS	Enable/Disable

- Prelude IDS

Prelude is an Hybrid IDS, combining "network intrusion detection" and "host based intrusion detection".

- Snort IDS

Snort is an open source Network IDS.

5.6. Services Activation

Services Summary

Refresh

	Status	
usb	Details--->	reload restart stop start Details
iptables	Details--->	reload restart stop start Details
harddrake	Details--->	reload restart stop start Details
kudzu	Details--->	reload restart stop start Details
ipvsadm	Details--->	reload restart stop start Details
network	Details--->	reload restart stop start Details
portmap	Running	reload restart stop start Details
syslog	Running	reload restart stop start Details
nfslock	Running	reload restart stop start Details
random	Details--->	reload restart stop start Details
netfs	Details--->	reload restart stop start Details
shorewall	Details--->	reload restart stop start Details
atd	Running	reload restart stop start Details
named	Running	reload restart stop start Details
ntpd	Stopped	reload restart stop start Details
sshd	Running	reload restart stop start Details

This page lists the services present on your machine. You will be given the opportunity to enable or disable them.

```
Status
[...]
Gpm Running reload restart stop start Details
Httpd-naat Running reload restart stop start Details
Squid Running reload restart stop start Details
[...]
```

The first column of the table lists the name of the service and its present status:

- **Running:** The service is installed and accepting connections.
- **Stopped:** The service is installed on the firewall but is currently disabled.
- **Unknown:** For some reason the Interface was unable to determine the status of that service.
- **Details--->:** Click on the "Details" button at the end of the row to get more information.

The parameters of these services may then be modified:

- **reload:** allows the reloading of that service configuration without interruption. To be used when a parameter of that service has just been modified.
- **restart:** stops and restarts the service.
- **stop:** the service will refuse further connections and terminate current ones.
- **start:** the service will accept further connections.
- **Details:** brings up another page with more information about that particular service.

Chapter 6. Configuring The Actual Firewall Behavior

In this chapter, we will go through all of the configuration pages of the interface's Firewall Rules section. The latter enables you to allow or deny traffic between the different zones and computers the firewall deals with.

6.1. Firewall Main Control



This section of the interface allows to control all the traffic that comes in and goes out from the firewall machine. It particularly allows to define the different groups of computers (zones) your firewall will deal with, and the traffic allowed between those zones.


In this introductory screen you can find the main control commands for all those filtering and routing services. There are four actions that can be taken on the whole firewall:

- start: starts the firewall and all associated services defined in this section;
- stop: stops the firewall. All communication channels will be closed, completely isolating the machine from the outside. That may be useful when you realize the machine is being compromised, but then the administration interface won't be accessible anymore but connecting physically from the firewall console.
- restart: stops the firewall (if it's running) and then starts it again;
- clear: erases all the configuration made by the administrator in the whole "firewall rules" section. Use with care! It then reinstalls and activate the default (factory) configuration. Use "start" or "restart" to come back to the custom configuration.



Be particularly aware of the implication when you use the "stop" or "clear" actions. When you have chosen the desired action, click on the "Next" button to confirm it.

6.2. Zones Definition



This sub-section allows to finely define each of the groups of computers (zones) the firewall will have to deal with. The introductory screen sums-up the current configuration and allows to manage the three components of the zones definition process.



Do not forget to click the "Apply" button when you have finished configuring the zones in this sub-section.

First of you have to define the zone names. Think carefully of the zones that may be necessary to your current network configuration. Three names are provided by default, as you can see in the first table which sums up the defined names. This default allows a simple yet safe configuration of your network.

- LAN: Local Area Network. The internal network where are connected internal client machines. Connections from the outside of this zone to machines of the LAN network are generally not allowed.
- DMZ: DeMilitarized Zone. Generally reserved for Internet servers. This zone will host computers dedicated to offer services to the Internet (the WAN zone defined below). Therefore connections to computers of this zone are allowed.
- WAN: Wide Area Network. This generally designates the Internet. Or more generally a network linked to the Internet.

For each of the defined zone names, click on the corresponding  icon to modify the names associated to that zone or  to definitely remove that zone.

If you wish to define a new zone, click on the "Add Zone"  icon.



There is a special zone "fw" not listed here but that always exists. It is used to designate the "firewall" zone: a zone made of one single machine, the firewall server itself.

Then it is necessary to inform the system of each network interface configured on your firewall, and the zone associated to them.

The table here lists the interfaces and the associated zones. If the zone name is "-" that means that various zones are attached to this interface. Those special "host" zones are defined in the third part of the page, below.

For each of the defined interfaces, click on the corresponding  icon to modify the zone associated to that interface or the options, or  to definitely remove that interface.

If you wish to add a new interface, click on the "Add Interface"  icon.

Finally, you can define in the last part of the page the possible "host" zones. Those zones are made of a group of computers sharing a single Ethernet interface of the firewall with other computers. For some reason you want to separate the way those machines are treated with respect to the other machines connected to that same interface. The machines owned by a "host" zone are identified by their subnet mask.

For example, this might be useful if your Internet server is physically connected to your LAN. You simply have to associate the DMZ zone to a "host" zone made of a single machine: the Internet server.

For each of the defined "host" zone, click on the corresponding  icon to modify its configuration, or  to definitely remove it.

If you wish to add a new "host" zone, click on the "Add Host"  icon.

6.2.1. Editing a Zone Identification



This form allows to add/modify a zone identification names. there are three strings identifying a single zone, used depending on the place it is displayed.

Zone ID (unique):	This ID number will be used everywhere needed to uniquely identify the zone. It is recommended not to modify the proposed default value.
Zone:	short name for the zone. The name should be 5 characters or less in length and consist of lower-case letters or numbers. It must begin with a letter. Also, the name assigned to the firewall and "multi" are reserved for system use.
Zone Display:	The zone name as displayed in logs.
Comments:	A string to identify more precisely the role of that zone, for maintenance purposes.


Warning1: No space characters are allowed in any of those three strings. It is safer to limit the characters used to letters, numbers, and underscore "_".

Warning2: If you rename or delete a zone, you should perform a "stop" followed by a "start" to install the changes rather than "restart" in the "Firewall Rules" section's main page.

Example: you have a farm of Web servers all located on the same subnet. We will create here that zone and configure it later.

Zone:	www
Zone Display:	WWW
Comments:	Web_Servers_Farm

6.2.2. Associating Zones to Interfaces



To each Ethernet interface that counts the firewall must be associated at least one zone. Associating multiple zones to a single interface is made possible through "host" zones. This form also allows to finely configure the options associated to the interface.

Interface ID:	This ID number will be used everywhere needed to uniquely identify the interface. It is recommended not to modify the proposed default value.
Zone:	Choose the zone you want to associate with the interface in the pull-down list. The special zone "-" means that various "host" zones will be associated to that interface.
Interface:	Choose the interface you want to configure in the pull-down list. If the desired interface is not shown, you need to declare it first in the "System setup" section.
Broadcast:	The broadcast address for the sub-network attached to the interface. This should be left empty for P-T-P interfaces (ppp*, ipp*); if you need to specify options for such an interface, enter "-" in this column. If you supply the special value "detect" in this column, the firewall will automatically determine the broadcast address.
Options:	Nine checkable options to specialize the interface behavior. See the table below.

Below are details about each of the options available for the interfaces. Review them all carefully for each interface, for some particular interfaces, some options are highly recommended.


dhcp	The interface is assigned an IP address via DHCP or is used by a DHCP server running on the firewall. The firewall will be configured to allow DHCP traffic to and from the interface even when the firewall is stopped.
-------------	--

noping	ICMP echo-request (ping) packets will be ignored by this interface.
routestopped	When the firewall is stopped, traffic to and from this interface will be accepted and routing will occur between this interface and other routestopped interfaces.
norfc1918	Packets arriving on this interface and that have a source or destination address that is reserved in RFC 1918 (Private network addresses) will be logged and dropped. This option is generally used for Internet Interfaces.
routefilter	Invoke the Kernel's route filtering facility on this interface. The kernel will reject any packets incoming on this interface that have a source address that would be routed outbound through another interface on the firewall. Warning: If you specify this option for an interface then the interface must be up prior to starting the firewall.
multi	The interface has multiple addresses and you want to be able to route between them. Example: you have two addresses on your single local interface eth1, one each in subnets 192.168.1.0/24 and 192.168.2.0/24 and you want to route between these subnets. Because you only have one interface in the local zone, Shorewall won't normally create a rule to forward packets from eth1 to eth1. Adding "multi" to the entry for eth1 will cause Shorewall to create the loc2loc chain and the appropriate forwarding rule. It is recommended to choose this option for the ppp interfaces.
dropunclean	Packets from this interface that are selected by the 'unclean' match target in iptables will be optionally logged and then dropped.
logunclean	This option works like dropunclean with the exception that packets selected by the 'unclean' match target in iptables are logged but not dropped.
blacklist	This option causes incoming packets on this interface to be checked against the blacklist. See the "blacklist" sub-section.

Example: with the same example of web servers farm, we will indicate now that the zone "www" is attributed the subnetwork connected on interface "eth3".

Zone:	www
Interface:	eth3
Broadcast:	detect

6.2.3. Associating "Host" Zones to Interfaces



If you have assigned the special zone "-" to an interface, you need now to define the "host" zones for that interface. A "host" zone is merely a group of machines identified by their common subnet. It can be reduced to a single machine.

Host ID:	This ID number will be used everywhere needed to uniquely identify the host zone.
Zone:	Choose the zone name to use for this host zone in the pull-down list.
Interface:	Choose the interface associated to this host zone in the pull-down list. Choose "+" if you don't want to associate a particular interface to the zone address or subnet. Note: Use of "+" weakens the firewall slightly and increases packet latency slightly.
IP Address:	The host or subnet address for the machines associated to this host zone. Example: "192.168.2.0/2".
options:	The "routestopped" option, if checked, has the following effect: When the firewall is stopped, traffic to and from this host (these hosts) will be accepted and routing will occur between this host and other routestopped interfaces and hosts.

Example: you wish to have some precise machines of the local network to be able to administer the firewall, even if the firewall is stopped. After having configured the local "eth2" interface as being assigned the special zone "-" with option "multi", you now need to configure the special host zone "adm" corresponding to only some machines of the local subnetwork.

Zone:	adm
Interface:	eth2
IP Address:	192.168.1.0/25
options:	routestopped

6.3. Masquerading, Static NAT and ProxyARP Configuration



This sub-section proposes services to ease the communications between the Internet and the internal machines, either clients or servers.

Warning: before leaving, do not forget to click the "Apply" button when you have finished configuring the needed services of this sub-section.

The "Classical Masquerading" part manages the rules allowing internal clients to access the Internet.

For each of the defined rule (if any), click on the corresponding  icon to modify the rule or  to definitely remove that masquerading.

If you wish to define a new masquerading, click on the "Add Masq"  icon.


The "Static NAT" part manages the rules for Network Address Translation. This allows internal servers (generally in the DMZ) to appear as being part of the Internet for external clients.

For each of the defined rule (if any), click on the corresponding  icon to modify the rule, or  to definitely remove that NAT.

If you wish to add a new NAT rule, click on the "Add NAT"  icon.

Finally, you can define in the last part of the page "Proxy ARP rules".

6.3.1. Classical IP Masquerading



To allow the clients of you internal network to access the Internet, you need to masquerade this network with respect to the Internet, as it is based on private addresses invalid on the Internet.

ID:	The unique ID number identifying this classical IP masquerading rule.
Masqueraded Network:	The subnet that you want to have masqueraded through the interface below. This may be expressed as a single IP address, a subnet or an interface name. The subnet may be optionally followed by "!" and a comma-separated list of addresses and/or subnets that are to be excluded from masquerading.
Through Interface:	The interface that will masquerade the subnet; this is normally your Internet interface.
Optional Network/Host:	You can optionally specialize the rule by adding a subnet or host IP. When this qualification is added, only packets addressed to that host or subnet will be masqueraded.
Source Address (SNAT) optional:	The source address to be used for outgoing packets. This column is optional and if left blank, the current primary IP address of the interface is used.

Example1: You have a number of IPSEC tunnels through ipsec0 and you want to masquerade traffic from your 192.168.9.0/24 subnet to the remote subnet 10.1.0.0/16 only.

Masqueraded Network:	192.168.9.0/24
Through Interface:	ipsec0
Optional Network/Host:	10.1.0.0/16
Source Address (SNAT):	

Example2: You have a DSL line connected on eth0 and a local network (192.168.10.0/24) connected to eth1. You want all local->net connections to use source address 206.124.146.176. Furthermore, you wish to exclude 192.168.10.44 and 192.168.10.45 from the SNAT rule

Masqueraded Network:	192.168.10.0/24!192.168.10.44,192.168.10.45
Through Interface:	eth0

Optional Network/Host:	
Source Address (SNAT):	206.124.146.176

6.3.2. Static Network Address Translation



Static NAT is a way to make systems behind a firewall and configured with private IP addresses (those reserved for private use in RFC1918) appear to have public IP addresses. To allow the clients of your internal network to access the Internet, you need to masquerade this network with respect to the Internet, as it is based on private addresses invalid on the Internet.

IMPORTANT: If all you want to do is forward ports to servers behind your firewall, you do NOT want to use static NAT. Port forwarding can be accomplished with simple entries in the "rules" sub-section. Also, in most cases Proxy ARP provides a superior solution to static NAT because the internal systems are accessed using the same IP address internally and externally.

ID:	The unique ID number identifying this static NAT rule.
External Public IP:	External IP address for the translation - This should NOT be the primary IP address of the interface named in the next field.
On this Network Interface:	Interface that you want the "External Public IP" address to appear on.
Internal Private IP (RFC1918):	Internal IP address for the translation. It must be a private IP address as defined by RFC 1918.

Two options are available for the translation:

All Hosts	If activated, this NAT will be effective from all hosts. If not then NAT will be effective only through the interface named in the "On this Network Interface" field.
Firewall system	If activated, the NAT will be effective also from the firewall system itself.

Example: We want to make the internal system with IP 10.1.1.2 appear to be on the Internet 130.252.100.* subnet. If we assume that the interface to the Internet is eth0, then the following rule would make the 10.1.1.2 system appear to have IP address 130.252.100.18.

External Public IP:	130.252.100.18
On this Network Interface:	eth0
Internal Private IP:	10.1.1.2
All Hosts	Yes
Firewall system	Yes

Note 1: The "All Hosts" option is used to specify that access to the external IP from all firewall interfaces should undergo NAT. If set to "No", only access from the interface in the "Interface" field should undergo NAT.

Note 2: Setting the The "Firewall system" option makes the packet originating on the firewall itself and destined for the "External" address to be redirected to the "Internal Private IP".

6.3.3. Proxy ARP's rule



This form is used to define proxy ARP (Address Resolution Protocol) rules. You need one rule for each system to be proxy ARP'd.

ID:	The unique ID number identifying this Proxy ARP's rule.
Server IP Address:	Address of the target system.
Internal Interface:	The interface that connects to the system. If the interface is obvious from the sub-netting, you may choose "-".
External Interface:	The external interface that you want to honor ARP requests for the "Server IP Address" specified above.
Have already a Route to Server IP:	If you already have a route through the "Internal Interface" to the "Server IP Address", Check this option. If you want the firewall itself to add the route, make sure it is not checked.

Example: You have public IP addresses 155.182.235.0/28. You configure your firewall as follows:

eth0 - 155.186.235.1 (Internet connection)
eth1 - 192.168.9.0/24 (masqueraded local systems)
eth2 - 192.168.10.1 (interface to your DMZ)

In your DMZ, you want to install a Web/FTP server with public address 155.186.235.4. On the Web server, you subnet just like the firewall's eth0 and you configure 155.186.235.1 as the default gateway:

Server IP Address:	155.186.235.4
Internal Interface:	eth2
External Interface:	eth0
Have already a Route to Server IP:	No

Note: You may want to configure the servers in your DMZ with a subnet that is smaller than the subnet of your Internet interface. In this case you will want to place "Yes" in the HAVEROUTE column.

6.4. Default Policies Configuration



This sub-section is used to describe the firewall policy regarding establishment of connections. Connection establishment is described in terms of clients who initiate connections and servers who receive those connection requests. Policies defined here are default policies. If no rule in the following "Rules" sub-section applies to a particular connection request then the default policy defined here is applied.

The table summarizes all the default policies currently configured. The factory settings defaults all policies to "REJECT", so that only connections explicitly allowed in the "Rules" sub-section are allowed.

Warning: Order is important: the firewall processes the policy rules from top to bottom and uses the first applicable policy that it finds. For example, in the following policy file, the policy for (lan, lan) connections would be ACCEPT as specified in the first entry even though the third entry in the file specifies REJECT.

If there are many rules, you can filter them by client and server zones. Choose the desired "client" and "Server"


zones available in the pull-down lists and click the  icon. The special "*" zone is simply a wildcard matching all zones.

Reminder: the "fw" zone designates the firewall itself.

For each of the defined policies, click on the corresponding  icon to modify that policy or  to definitely remove it. To add a new default policy, click the  icon.

Do not forget to click on the Apply button at the end of the page once all default policies are set in order to make your changes effective.

6.4.1. Defining a Default Policy



You are about to define here the default policy for a connection request between a client and a server zone.

For this policy to be activated, the connection must be originated by a machine from the "Client Zone" and directed to a machine belonging to the "Server Zone". Then the "Default policy" action will be taken for that connection. Optionally, if this policy has been activated, it will generate a log entry with level "Log Level".

Policy ID	The unique ID number identifying this policy rule.
Client Zone	The zone from which the connection must be originated for the policy to activate.
Server Zone	The zone to which the connection is targeted.
Default policy	The action that will be taken if the policy is actually activated. See the table below for details about all possible actions.
Log Level	If set to "-", no log message is generated when the policy is applied. Otherwise, a syslog message is generated with the indicated syslog level. See the syslog.conf man page for a description of each log level.

Here is a short description of the four possible policies:

ACCEPT	The connection is allowed.
DROP	The connection request is ignored.
REJECT	The connection request is blocked and a "destination-unreachable" message is sent back to the client.
CONTINUE	The connection is neither ACCEPTed, DROPped nor REJECTed. CONTINUE may be used when one or both of the zones named in the entry are sub-zones of or intersect with another zone.

Example: you trust the people that are on your local network "lan" and do not wish to restrict them access to any service on the Web (zone "wan"). You do not wish to log their activity either.

Client Zone	lan
Server Zone	wan
Default policy	ACCEPT
Log Level	-

6.5. Firewall Rules Configuration

MandrakeSecurity

Rules >

Rules Configuration

Rules configuration is used to express firewall rules that are exceptions to the high-level policies established in Default Policies

View Rules filtered on:

Client: * Server: * Port: *

Add Simple Rule **Add Custom Rule**

	Result	Client Zone	Server Zone	Protocol	Port(s)	Forward
1	ACCEPT	fw	wan	tcp+udp	53	
2	ACCEPT	dmz	wan	udp	53	
3	ACCEPT	lan	wan	udp	53	
4	REJECT	wan	fw	tcp	113	
5	ACCEPT	lan	fw	tcp	22	
6	ACCEPT	lan	fw	tcp	8443	
7	ACCEPT	fw	lan	icmp	8	
8	ACCEPT	lan	fw	icmp	8	
9	ACCEPT	lan	dmz	icmp	8	
10	ACCEPT	dmz	lan	icmp	8	
11	ACCEPT	dmz	fw	icmp	8	
12	ACCEPT	fw	dmz	icmp	8	
13	ACCEPT	lan	wan	tcp	pop3	
14	ACCEPT	lan	wan	tcp	smtp	
15	ACCEPT	lan	wan	tcp	http	
16	ACCEPT	lan	wan	tcp	https	
17	ACCEPT	lan	wan	tcp	ssh	

System Setup

Internet Access

Services

Firewall Rules

- Zones Setup
- Masq NAT
- Default Policies
- **Rules**
- BlackList
- TOS
- Tunnels

VPN

Monitoring

Tools

Logout

We are here in the very core of the firewall. The "Rules" sub-section defines exceptions to the policies established in "Default Policies". There is one entry in the table for each of these rules.

The table summarizes all the rules currently configured. By default, MandrakeSecurity define standard rules for the default zones (LAN, WAN, DMZ). As the default policy is to "REJECT" any connection, the default rules allow some precise ones:

- Computers from the LAN can connect to the Internet (WAN) for Web browsing, mail services, FTP, and SSH connections;
- Computers from the LAN can connect to the firewall's SSH server or MandrakeSecurity Web Interface;
- All DNS (Domain Name Service) requests towards the Internet are accepted;
- Ping echos are allowed between internal zones.

If there are many rules, you can filter them. Choose the desired "Client" and "Server" zones as well as a "Port"

in the pull-down lists and click the  icon. The special "*" zone or port is simply a wildcard matching all possibilities.

Reminder: the "fw" zone designates the firewall itself.

For each of the defined rules of the table, click on the corresponding  icon to modify that rule or

 to definitely remove it.


If you wish to add a new rule, two forms are actually available: click on the corresponding "Add Rule"

 icon:

Add Simple Rule	You will get here the simple rule form, allowing to define an "ACCEPT" rule only specifying the source, destination and protocol.
Add Custom Rule	The form displayed here allows the definition of more complex rules, with all types of actions available, and some options, such as logging, forwarding and SNAT.

Do not forget to click on the Apply button at the end of the page once all default policies are set in order to make your changes effective.

6.5.1. Defining a Simple "ACCEPT" Firewall Rule



You are about to define here a new rule to authorize a specific connection between two different zones.

Whenever a connection matches the criterion defined here, it will be allowed.

Rule ID	The unique ID number identifying this policy rule.
Pre-defined Services	Choose either a common service in the pull-down list, or enter a name or service number in the field.
Protocol	The protocol type associated to that service.
Coming from	The zone from which the connection request is originated.
and going to	the zone to which the connection request is directed, optionally followed by an IP or a subnet. Leave "-" in the field for the whole zone.
Forward	When this option is checked, the rule behavior is modified. All requests from the specified source and for the specified service will be caught, whatever the target system is. Then this request will be forwarded to the "and going to" IP. In this case, the "and going to" field must specify a specific IP address.

Example: You wish to forward all ssh connection requests from the Internet to local system 192.168.1.3.

Pre-defined Services	Secure remote connection [ssh]
Protocol	tcp
Coming from	wan
and going to	lan 192.168.1.3
Forward	[checked]

6.5.2. Defining a Complex Firewall Rule



You are about to define here a new rule to manage a specific connection between two different zones. If the request matches the different criterion defined here, the "Result" action will be taken.

Here is a description of the different fields available in the form, fill them according to the criterion you want to be matched for this rule to be activated. Some options are also available to manage these connections:

Rule ID	The unique ID number identifying this policy rule.
Result	The action taken for the connection request matching this rule. See table below.
Logging	Set to "info" if you want each of these connections logged by syslog when accepted.
Pre-defined Services	Choose either a common service in the pull-down list, or enter a name or service number in the field.
Protocol	The protocol type associated to that service.
Client	The zone from which the connection request is originated. The matching can be narrowed by specifying a precise IP or subnet, or even a port number. Leave "-" in the field for matching any IP or port.
Server	the zone to which the connection request is directed. The matching can be narrowed by specifying a precise IP or subnet, or even a port number. Leave "-" in the field for matching any IP or port.
Forwarding Address	If the request is targeted at the IP specified here (or if it is set to "all"), it will be forwarded to the "Server" IP and port. In this case, the "Server" field must specify a specific IP address.
SNAT	If specified, and if the forwarding is activated above, then the source address of the request will be set to this "SNAT" value before being forwarded to the server.

Here is a short description of the four possible actions:


ACCEPT	The connection is allowed.
DROP	The connection request is ignored.

REJECT	The connection request is blocked and a "destination-unreachable" message is sent back to the client.
CONTINUE	The connection is neither ACCEPTed, DROPped nor REJECTed. CONTINUE may be used when one or both of the zones named in the entry are sub-zones of, or intersect with, another zone.

Example: you want the FTP server on 192.168.2.2 in your masqueraded DMZ to be accessible from the local 192.168.1.0/24 subnetwork. Note that since the server is in the 192.168.2.0/24 subnetwork, we can assume that access to the server from that subnet will not involve the firewall.

Result	ACCEPT
Logging	
Pre-defined Services	ftp
Protocol	tcp
Client	lan 192.168.1.0/24
Server	dmz 192.168.2.2
Forwarding Address	155.186.235.151
SNAT	

6.6. Maintaining the BlackList



This page lists the hosts or subnets for which connections requests will be systematically dropped, even if an explicit rule allows the connection normally. Note that for that list to be effective, the interface(s) to which those IPs are connected must have the "blacklist" options set in the "Zones setup" sub-section.

- To add a new blacklisted IP/subnet, enter it in the "Add Host IP/Subnet" field and click the "Add Entry"



icon. The new address will appear in the list.

- To remove a blacklisted IP/subnet, select it in the list and click the "Delete Entry"



icon.

Do not forget to click on the Apply button at the end of the page once you are satisfied with your settings in order to make your changes effective.

6.7. Type Of Service Rules Configuration

TOS >

TOS Rules Configuration

TOS Rules configuration define the Type of Service field in packet headers based on packet source, packet destination, protocol, source port and destination port.

View Rules filtered on:

Client: all Server: all Port: ssh

Add TOS Rule

	Client Zone	Server Zone	Protocol	Client Port	Server Port	TOS
1	all	all	tcp		ssh	16
2	all	all	tcp	ssh		16
3	all	all	tcp		ftp	16
4	all	all	tcp	ftp		16
5	all	all	tcp		ftp-data	8
6	all	all	tcp	ftp-data		8

Buttons: ? Help, X Cancel, Apply

This page lists and allow to manage the TOS (Type Of Service) rules of the firewall. TOS rules direct the firewall to modify packet headers by adding a TOS value. This value gives additional information to routers notably, so that the packets are given the optimal treatment regarding their use.

The table summarizes all the TOS rules currently configured.

If there are many rules, you can filter them. Choose the desired "Client" and "Server" zones as well as a "Port" in the pull-down lists and click the  icon. The special "*" zone or port is simply a wildcard matching all possibilities.


Reminder: the "fw" zone designates the firewall itself.

For each of the defined rules of the table, click on the corresponding  icon to modify that rule or  to definitely remove it.

If you wish to add a new rule, click on the "Add TOS Rule"  icon.

Do not forget to click on the Apply button at the end of the page once you are satisfied with your settings in order to make your changes effective.

6.7.1. Editing TOS Rules



The screenshot shows the 'Add TOS Rule' window in the MandrakeSecurity interface. On the left, a sidebar lists various configuration options, with 'Firewall Rules' expanded and 'TOS' highlighted. The main window contains a form with the following fields: 'Rule Id' (text box with '7'), 'TOS' (dropdown menu with 'Maximize Throughput (8)'), 'Zone' (text box with 'Ex: Zone:[Interface,IP or Subnet]:[Port]'), 'Client' (dropdown menu with 'all'), 'Server' (dropdown menu with 'all'), and 'Protocol' (dropdown menu with 'tcp'). At the bottom of the form are four buttons: '? Help', 'X Cancel', 'Back <', and '> Next'.

This form defines TOS (Type Of Service) rules, which add a TOS value to matching packet headers.

Whenever a packet passing through the firewall matches the criterion defined here, it will be added the corresponding TOS value.

Rule ID	The unique ID number identifying this TOS rule.
TOS	Choose the type of treatment that will best suite the kind of packets matched by this rule.
Client	The zone from which the packet is originated. The matching can be narrowed by specifying: a precise IP, interface or subnet in the middle field; or even a port number in the right field. Leave "-" in a field for matching any IP or port.
Server	the zone to which the packet is directed. The matching can be narrowed by specifying a precise IP, interface or subnet in the middle field; or even a port number in the right field. Leave "-" in a field for matching any IP or port.

Protocol	The protocol type associated to that service.
-----------------	---

Example: you want FTP data packets to be handled so that the rate transfer is maximized, disregarding reliability and delay, in all directions.

TOS	Maximize-Throughput (8)
Client	all
Server	all
Protocol	ftp-data

Chapter 7. VPN Configuration

This chapter explains what a VPN is and how to configure a VPN using *MandrakeSecurity* to act both as a server and client side of the VPN.

7.1. What is a VPN?

VPN stands for Virtual Private Networking and is a way to establish private networks over public networks, like the Internet, in a secure way. See figure 7-1.




Figure 7-1. VPN Connection Scheme

VPN uses tunneling to create a private network that goes across the public Internet. Tunneling can be thought of as the process which consists of encapsulating a packet inside another one and sending it over the network (in VPN's case, the Internet). *MandrakeSecurity* uses IPSec¹ as the tunneling protocol.

The “network of trust” between private machines across the public Internet is built upon certificates and a Certificate Authority (CA). The CA is an entity **trusted** by the VPN participants. The certificates created by your *MandrakeSecurity* system adhere to industry standards but they will not be guaranteed by a third party CA such as VeriSign, for example. Of course, you can also use your own certificates approved by public CAs with *MandrakeSecurity*.

7.2. Why a VPN?

Setting a VPN up has many advantages among which are the following:

- Ability to link together geographically-distributed private networks. This is the very reason why VPNs were developed in the first place. Think about linking the different subsidiaries of your company together no matter where they are in the world.
- Secure communications. Since all VPN traffic is encrypted you can rest assured that all your data is safe while traveling over the network.
- Costs reduction. By using an already established across-the-world network (the Internet) all you need to interconnect your distributed private networks already exists. No need to pay for (very) expensive dedicated communications links. Using only a few high speed Internet links like DSL ones is more than enough.

1. IPIP is also supported.



Actually, VPNs can also be set up across analog modem links, but usually a permanent type of connection (like DSL or cable-modem) is preferred. This will depend on the way you use the VPN.

Of course, there are also (few) drawbacks:

- You do not manage the whole network. Since public networks are an inherent part of a VPN network, the “public” part of the network is totally out of your control.

However, this is also one of its greatest advantages: less network management and reduced costs. Besides, the public part of the network is the Internet and the latter is well managed by knowledgeable people with enough resources to assure the needed service quality.

- Single point of failure. Usually, there is only one VPN server so, if that machine fails, the network goes down. However this risk can be diminished by using fault-tolerant machines. As always, it all depends on how critical your operations are.

7.3. Setting up a VPN Server

The following sections will detail a VPN server setup on a *MandrakeSecurity* system. It is assumed that the *MandrakeSecurity* system has a permanent connection to the Internet with a fixed IP address and a LAN behind it with a 192.168.0/24 subnetwork. Also, the *MandrakeSecurity* system will act as the CA for the VPN.

The configuration steps will be presented in a logical order. It is **highly** recommended that you follow them in that order. However, if you know what you are doing, you might perform some of the steps as you wish.

7.3.1. Back up MandrakeSecurity's Current Configuration

Since a VPN setup involves changes to the firewall policies and rules, it is a good idea to back those up using the Backup and Restore facility provided in the Tools section. Please refer to “*Management Tools*”, page 127 for more information.

7.3.2. Create a VPN Zone

Go to the Zones Setup sub-section of the Firewall Rules section and click on the Add Zone link. Example values for the fields are shown in the figure below:

Zone Add

This form allows to add/modify a zone identification names. there are three strings identifying a single zone, used depending on the place it is displayed.

Zone ID (unique):	5
Zone:	vpn
Zone Display:	VPN
Comments:	vpn_zone

? Help X Cancel Back < > Next

Figure 7-2. Adding a VPN Zone

Please refer to “*Configuring The Actual Firewall Behavior*”, page 67 for more information about the meaning about the different fields.



No space characters are allowed in any of those strings. It is safer to limit the characters used to letters, numbers, and the underscore (_) character.

Then press the Next button to add the new zone identifying the VPN itself. Once you see the zone listed in the Zones Setup page, press the Apply button to make your changes effective.

7.3.3. Add the VPN Network Interface

Go to the Zones Setup sub-section of the Firewall Rules section and click on the Add Interface link. Example values for the fields are shown in the figure below:

Figure 7-3. Adding an ipsec Network Interface

Please refer to “*Configuring The Actual Firewall Behavior*”, page 67 for more information about the meaning about the different fields.

Then press the Next button to add the IPSec interface for the VPN. Once you see the interface listed in the Zones Setup page, press the Apply button to make your changes effective.




ipsecN interfaces are logical interfaces bound to a physical interface, such as eth0 for example. You can associate more than one VPN with the same ipsecN interface and add different VPN zones to set up per-zone policies, if you need to.

7.3.4. Add Firewall Default Policies for VPN Traffic

Now you need to add default firewall policies to handle VPN traffic. To do so, go to the Default Policies sub-section of the Firewall Rules section and click on the Add Policy link. The figure below shows the policy for

traffic coming in from any zone and going towards the VPN (all->vpn):



Add Policy

Policy Id	1
Client Zone	all
Server Zone	vpn
Default Policy	ACCEPT
Log Level (optional)	info

? Help X Cancel Back Next

Figure 7-4. Adding Default Policies for the VPN

You also need to add a similar policy for traffic coming in from the VPN and going towards all the other zones (vpn->all) in order to set up a bidirectional communications link.

Please refer to *“Configuring The Actual Firewall Behavior”*, page 67 for more information on the meaning about the different fields.

Then press the Next button to add the firewall policy for the VPN. Once you have added the two policies and see them listed in the Default Policies Configuration page, press the Apply button to make your changes effective.




This setup will allow **all and any** kind of VPN traffic. It is OK to leave it as is for initial VPN setup and testing, but once you are sure the VPN is established and working, you should change the two policies defined above to the REJECT behavior and add rules for the specific kinds of traffic you want to allow over the VPN. Please refer to *Testing the VPN and Making it More Secure*, page 98 for an example on how to set up rules to allow HTTP traffic.

7.3.5. Add the Firewall's VPN Tunnel

It is time to add the VPN “tunnel” in the firewall, allowing traffic on port 500/udp. Go to the Tunnels sub-section of the Firewall Rules section and click on the Add Tunnel link. The figure below shows values for the

tunnel:



Tunnel Add

ID (unique):	1
Type:	ipsec
Zone:	wan
Gateway IP:	0.0.0.0/0
Gateway Zone(optional):	vpn

? Help X Cancel Back ← → Next

Figure 7-5. Adding a Firewall Tunnel

- **ID (unique).** The unique identifier for this tunnel. It is highly recommended to leave this value unchanged.
- **Type.** The tunnel type. `ipsec` for an IPSec tunnel (the default and recommended setting); `ipip` for an IPIP tunnel.
- **Zone.** The zone from/to where VPN traffic will flow using this tunnel. For the kind of VPN we are setting here, this must be set to `wan` (the Internet zone).
- **Gateway IP.** The IP address of the “remote” gateway machine. In our example we set this to `0.0.0.0/0`, meaning that VPN traffic will be allowed from anywhere on the Internet.



Setting this to `0.0.0.0/0` can be thought of as a security risk and indeed it is not very secure. However it is the **only** possible setting for machines with non-permanent Internet connection types (like analog modems) and machines with non-fixed IP addresses (almost all DSL and cable-modem connections).

- **Gateway Zone (optional).** Set it to `vpn` because the VPN will be the “gateway” between both private networks.

Then press the Next button to add the firewall tunnel for the VPN. Once you see it listed in the Tunnels page, press the Apply button to make your changes effective.

7.3.6. Generate the CA Certificates

All parties involved in a VPN need a certificate as a “proof” of their identity (authentication) and for encryption purposes.



The following is only needed if your *MandrakeSecurity* is going to be the Certificate Authority (CA) of your VPN. If it is not your case, you can safely ignore this step.



It is **mandatory** to create the CA Key first.

7.3.6.1. CA Key

Go to the CA sub-section of the VPN section and click on the CA Key link and press the Next button. The figure below shows example values for the fields:

Edit VPN CA Configuration	
Common Name	mandrake.com
Days	3650
Crl Days	33
Bits	2048
Country	fr
State or Province	fr
Locality	paris
Organization Name	mandrake
Organizational Unit Name	mandrake
Email Address	admin@mandrake.com

? Help X Cancel Back Next

Figure 7-6. Configuring the CA



No space characters are allowed in any of those strings. It is safer to limit the characters used to letters, numbers, and the underscore (_) character.

Following is a brief explanation of some of the fields (the others are self-explanatory):

- **Common Name.** This **must** be set to the FQDN (Fully Qualified Domain Name) of your *MandrakeSecurity* host.
- **Days.** The expiration time of this certificate in days. Set to 10 years in the example.
- **Crl Days.** How many days until the Certificate Revocation List (CRL) is considered obsolete.
- **Bits.** How many bits to use for the key generation. Normally set to 1024 or 2048. Do not use values less than 1024 for this field.
- **Country.** The two-letter ISO code for the country where your *MandrakeSecurity* system resides.
- **E-mail Address.** The e-mail address of the *MandrakeSecurity* system administrator. Normally, this is set to `admin@fqdn_of_the_mandrakesecurity_machine.ext`.

Once you are satisfied with your settings, press the Next button and then click on the Generate Autosigned Certificate link to generate the certificate for the CA.

7.3.6.2. Other Keys

Go to the CA sub-section of the VPN section and click on the Other Keys link and then on the Add VPN Entries link. You will see the same dialog as on the CA Key page.

It is **mandatory** for the first entry you add to be the VPN server one, that is your *MandrakeSecurity* system in our example. Now go ahead and enter the same values you did for the CA Key and click on the Next button.

When this is done, you need to add entries for the remote points. Once you have added all entries and see them listed in the VPN CA Other List Configuration page, press the Apply button to make your changes effective.

7.3.7. Add a VPN Server

The figure below shows the VPN we are discussing, including details on the private network's IP and mask and the meaning of the "left" and "right" sides needed in the configuration to be performed in this step:




Figure 7-7. VPN Layout

So, your *MandrakeSecurity* server (MNF Server A) is on the “left side” and every other server/client is on the “right side” of the VPN. This is a convention that must be established beforehand and both sides (left and right) must be set up for the configuration to be complete.

Go to the Server sub-section of the VPN section and click on the Add VPN Server link. The figure below shows typical values for the fields:



Add a VPN Server	
VPN Server Id	1
Side	Left
CommonName	mandrake.com
IP	public_IP_mandrake.com
Subnet/Netmask	192.168.0.0/24
Nexthop	gway_IP_mandrake.com
Authentication	x509

Buttons: ? Help, X Cancel, Back, Next

Figure 7-8. Adding a VPN Server Side



The VPN server **must** be the first entry.

- **VPN Server ID.** Numerical Unique Identifier. It is safe (and recommended) to leave this as is.
- **Side.** Set to left for your *MandrakeSecurity* system and to right for the remote system.
- **Common Name.** This **must** be set to the FQDN of your *MandrakeSecurity* host for the left side and to the FQDN of the remote host for the right side.
- **IP.** The IP address on the Internet interface of your *MandrakeSecurity* host for the left side and the IP address of the remote host for the right side.
- **Subnet/Netmask.** The network IP and mask for the corresponding (left or right) side. In our example, we set it to 192.168.0.0/24 for the left side and to 172.16.1.0/24 for the right side.
- **Nexthop.** The IP address of the system’s gateway. This will be dependent of the host’s IP address you are configuring, but usually it is the same as the host’s IP but with 1 as the last number. For example, if your host’s IP is 123.234.123.200, then this value should be set to 123.234.123.1.

- **Authentication.** x509 is the only certificate type supported.

Once you have added both left and right sides, click on the Apply button and then on the Restart IPSEC link to make your changes effective.

Congratulations! The VPN server is already set up.

7.3.8. Distributing the Certificates and Keys



The following is only needed if your *MandrakeSecurity* is going to be the Certificate Authority (CA) of your VPN. If it is not your case, you can safely ignore this step.

Now you have to distribute all needed certificates and keys to the interested parties. This distribution must be done in a **secure** way because the security of the whole VPN depends on those certificates and keys. You can mail the needed files using encryption (ex: with *OpenPGP*), meet the interested parties personally and hand them diskettes holding the data, or any other secure way you can think of. **Never** e-mail the files without using some form of encryption.

The files to distribute to the remote parties are:

- /etc/freeswan/ipsec.d/fqdn_of_mandrakesecurity.crt
- /etc/freeswan/ipsec.d/fqdn_of_remote_system.crt
- /etc/freeswan/ipsec.d/private/fqdn_of_remote_system.key

The remote parties must then copy those files to the appropriate places on their systems; needless to say, this is system-dependent so it will not be detailed here.



After the certificates and keys were distributed and copied to the appropriate places on the remote hosts, the IPSec service has to be restarted on those remote hosts in order to include the distributed certificates and keys.

7.3.9. Testing the VPN and Making it More Secure

It is now time to test the VPN. This is rather simple. Just ping the remote network and see if you get an answer. Following our example, the command `ping -c 2 172.16.1.10` should return something like this

```
PING 172.16.1.10 (172.16.1.10) from 192.168.0.70 : 56(84) bytes of data.  
64 bytes from 172.16.1.10: icmp_seq=1 ttl=64 time=0.047 ms  
64 bytes from 172.16.1.10: icmp_seq=2 ttl=64 time=0.069 ms  
--- 192.168.0.70 ping statistics ---  
2 packets transmitted, 2 received, 0% loss, time 1502ms  
rtt min/avg/max/mdev = 0.047/0.055/0.069/0.011 ms
```

assuming the remote host, 172.16.1.10, is up and running.

Once you are sure it is working you should replace the default firewall policies for the VPN to REJECT (deny all VPN traffic by default) and add firewall rules for the specific services you need to use over your VPN. The following figure shows the field values for a firewall rule allowing HTTP traffic on port 80 over the VPN:

Add Rule

Select pre-defined services or enter custom services allowed to enter in the Zone

Rule Id	33	
Services	Pre-defined Services or Custom Port(s)	
Pre-defined Services	Web, internet pages [http]	
Protocol	tcp	
Coming from:	wan	
and going to:	Ex: Zone:[Interface,IP or Subnet]	
	vpn	
Forward	☐	

? Help

X Cancel

Back

Next

Figure 7-9. Rule to Allow HTTP Traffic Over the VPN

7.4. Setting up a VPN Client

Go to the Client sub-section of the VPN section and click on the Next button. The figure below hints on the values to enter in the different fields:

Local		Remote	
Side	Left Empty list...	Side	Right Empty list...
CommonName	mandrake.com	CommonName	remote.com
IP	public_IP_mandrake.com	IP	public_IP_remote.com
Subnet	192.168.0.0/24	Subnet	172.16.1.0/24
Nexthop	gway_IP_mandrake.com	Nexthop	gway_IP_remote.com

? Help X Cancel Back < > Next

Figure 7-10. Adding a VPN Client Side

Once you have completed all fields press the Apply button to make your changes effective.



Please refer to *Add a VPN Server*, page 96 for more information about the fields' meaning.



Once the VPN is set up the client/server roles do not have to remain the same. For example, you could set up a VPN server on your network and make the services (servers) accessible through it on the remote network.

You need to copy the files listed in *Distributing the Certificates and Keys*, page 98 to your system (remember this is system-dependent and thus, we will not explain how to do it in this chapter) and then click on the Restart IPSEC link for the settings to be taken into account.

Chapter 8. Configuring Masqueraded Clients

This chapter will show you how to make different operating systems use a *GNU/Linux* box with masquerading set up as a gateway to the outside world. The configuration tests on the following operating systems all proved successful:

- Apple Macintosh, with MacTCP or Open Transport;
- Commodore Amiga, with AmiTCP or AS225-stack;
- Digital VAX Stations 3520 and 3100, with UCX (TCP/IP stack for VMS);
- Digital Alpha/AXP, with Linux/Redhat;
- IBM AIX (on RS/6000), OS/2 (including Warp 3) and OS400 (on OS/400);
- Linux (of course!): any kernel release since 1.2.x;
- Microsoft DOS (with the NCSA Telnet package, partial DOS Trumpet support), Windows 3.1 (with the Netmanage Chameleon package) and Windows For Workgroup 3.11 (with TCP/IP package);
- Microsoft Windows 95, Windows 95 OSR2, Windows 98, Windows 98se;
- Microsoft Windows NT 3.51, 4.0 and 2000 (both workstation and server);
- Novell Netware 4.01 Server, with the TCP/IP service;
- SCO OpenServer (v3.2.4.2 and 5);
- Sun Solaris 2.51, 2.6 and 7.

Let's go through the configuration of a few of them. If your system is not listed, a simple way to proceed is to "just tell the OS which machine to use as a gateway". Note that our main focus here is the **gateway** side of the network: therefore, we won't touch on DNS, file sharing or connection schemes problems. Thus, for this chapter to be of any use to you, you need a well-configured local network. Refer to your system's documentation to set it up properly, paying special attention to the DNS settings.

What follows assumes that you are set up on a class C network: your different machines all have IP addresses like 192.168.0.x, with a netmask set to 255.255.255.0, and use eth0 as the network interface. We also take for granted that your gateway's IP address is set to 192.168.0.1, and that your machines can each "talk" to the gateway (test the latter with the ping command or its equivalent in your environment).

8.1. Linux Box

There are (at least) three ways to go about this.

8.1.1. On-The-Fly Configuration

This is probably the fastest way to proceed. However, when you next restart your network layer or your whole system, any configuration change you made will have disappeared!

If eth0 is the network interface through which you access the gateway, (as root) issue this simple command: `route add default gw 192.168.0.1 eth0` That's it! If the gateway is properly configured and connected to the Internet, the whole world is now within your reach through your favorite web browser.

8.1.2. Permanent, Manual Configuration


To maintain the configuration each time the system is shut down and restarted, we need to edit a configuration file. Its name is `/etc/sysconfig/network` on a **Mandrake Linux** machine (it may be different on yours). Open it with your usual text editor, then add the following lines:

```
GATEWAYDEV="eth0" GATEWAY="192.168.0.1"
```

You may now restart your network layer with: `service network restart`

8.1.3. Permanent, Automatic Configuration


To install the configuration automatically, it's just about putting the right parameters in the configuration wizard. Refer to the *Starter Guide's Internet Configuration*. When you are configuring a local network Internet connection, the first step offers to configure the network in manual or automated mode (DHCP):

The image shows a window titled "Network & Internet Configuration" with a blue header bar containing a logo of a globe and a star. The main content area has a white background and contains the text "Configuring network device eth0 (driver 3c59x)". Below this, it says "Please enter the IP configuration for this machine. Each item should be entered as an IP address in dotted-decimal notation (for example, 1.2.3.4)." There are two input fields: "IP address" with the value "192.168.0.15" and "Netmask" with the value "255.255.255.0". Below these fields is a checkbox labeled "Automatic IP" which is unchecked, followed by the text "(bootp/dhcp)". At the bottom of the window are three buttons: "Cancel", "<- Previous", and "Next ->".

Configuring network device eth0 (driver 3c59x)	
Please enter the IP configuration for this machine. Each item should be entered as an IP address in dotted-decimal notation (for example, 1.2.3.4).	
IP address	192.168.0.15
Netmask	255.255.255.0
Automatic IP	☐ (bootp/dhcp)

Figure 8-1. Reconfiguring the Local Network with drakconnect

Simply put the right information in it. If you have a bootp or DHCP server on your local network, simply check the Automatic IP box, and your configuration is done. If you have a static IP address for your machine, enter it in the first field after making sure the Automatic IP check box is deactivated. Then click on the Next -> button.



The image shows a Windows-style dialog box titled "Network & Internet Configuration". It has a blue header bar with a globe icon on the left. The main area is white with a grey border. It contains a text box with instructions: "Please enter your host name. Your host name should be a fully-qualified host name, such as 'mybox.mylab.myco.com'. You may also enter the IP address of the gateway if you have one". Below this are three input fields: "Host name" with "test.mandrakesoft.com", "DNS server" with "192.168.0.10", and "Gateway (e.g. 192.168.0.1)" with "192.168.0.1". At the bottom are three buttons: "Cancel", "<- Previous", and "Next ->".

Field	Value
Host name	test.mandrakesoft.com
DNS server	192.168.0.10
Gateway (e.g. 192.168.0.1)	192.168.0.1

Figure 8-2. Setting up the Gateway with drakconnect

Here, you must write in the correct IP addresses for the gateway and DNS server. Once this is done, follow the wizard's steps and restart the network when proposed. And that's it. Your network is properly configured and ready to run. The configuration is now permanent.

8.2. Windows XP Box

We will assume here that you already have a configured network connection. The following snapshot shows the three different steps to get to the desired dialog.




Figure 8-3. Setting up The Gateway with Windows XP

Here are the actions to take to go from one window to another:

1. On the desktop, right-click on the My network places icon, and select Properties in the menu that appears.
2. In the Network Connections window, do the same with the connection linked to the network where the gateway is located.
3. In the next dialog, select the Internet Protocol (TCP/IP) entry and click the Properties button.
4. In this dialog, you can choose to check Obtain an IP address automatically if you have a DHCP server on your network. Then, the gateway should also be automatically configured. If not, check Use the following IP address and fill in the associated fields.

8.3. Windows 95 or Windows 98 Box



Figure 8-4. The Network Icon Under Windows 95

Start by going in the Control Panel (Start+Settings→Control Panel) and find the network icon as shown. Double-click on it: the network configuration panel comes up.




Figure 8-5. The Network Configuration Panel under Windows 95

In the displayed list, you should find a protocol named TCP/IP. If not, you will have to refer to your system documentation to find out how to install it. If it is already there, select it and click on “Properties”.




Figure 8-6. The TCP/IP Configuration Panel under Windows 95

This window will enable you to set up your TCP/IP parameters. Your system administrator will tell you if you have a static IP address or if you are using DHCP (automatic IP address). Click on the Gateway tab.




Figure 8-7. The Gateway Configuration Panel under Windows 95

The rest is child's play! Fill in the blanks with your gateway's IP address (i.e. 192.168.0.1, in our example). Click the Add then the OK buttons.

You will need to reboot your computer, of course. Once this is done, find out if you can reach the rest of the world.

8.4. Windows NT or Windows 2000 Box

To configure these OSs, follow these simple steps:

1. Go to Control Panel+Network→Protocol.




Figure 8-8. The Protocol Configuration Panel under Windows NT

2. First, select the TCP/IP Protocol in the list of network protocols. Then, click on the Properties button, and select the network card connected to the local network (figure 8-9). In this example, we show a configuration with the DHCP server activated on the *MandrakeSecurity* server: the Obtain an IP address from a DHCP server option is checked.




Figure 8-9. The Network Software Panel under Windows NT

If this is your case, you just need to confirm all those choices and reboot. Otherwise, follow the following steps.

3. If you have no DHCP server, you need to manually set all parameters. Begin by checking the Specify an IP address option (figure 8-10).




Figure 8-10. The TCP/IP Configuration Panel under Windows NT

Select the appropriate adapter, the IP address should already be correct.

4. Simply fill in the Default Gateway field with 192.168.0.1 (the address of the *Linux* box sharing the connection in our example).
5. Finally, you will need to specify the DNS servers you use in the DNS tab as shown in figure 8-11.




Figure 8-11. The DNS Configuration Panel under Windows NT

You must also provide a host name and an associated domain name.



Unless you know exactly what you are doing, proceed with utmost care with the following steps:

- leave the Automatic DHCP configuration field blank unless you have a DHCP server somewhere on your network;
- leave all the WINS Server fields blank as well unless you have one or more WINS servers;
- do not place a check in the Enable IP Forwarding field unless your NT machine is used for routing and, once again, you know perfectly what you are doing;
- please disable DNS for Windows Name Resolution and Enable LMHOSTS lookup.

Click on OK in the dialog boxes which then appear and restart your computer to test the configuration.

8.5. DOS Box Using the NCSA Telnet Package

In the directory which hosts the NCSA package, you will find a file called `config.tel`. Edit it with your favorite editor and add the following lines:

```
name=default host=yourlinuxhostname
hostip=192.168.0.1 gateway=1
```

Of course, write the name of your *Linux* box instead of `yourlinuxhostname` and change the gateway address given here (192.168.0.1), which is only an example.

Now save the file, try to telnet your *Linux* box, then a machine somewhere out there...

8.6. Windows for Workgroup 3.11

The TCP/IP 32b package should already be installed. Go to the Main+Windows Setup+Network Setup→Drivers menu entry and select Microsoft TCP/IP-32 3.11b in the Network Drivers section, then click Setup.

From here, the procedure is quite similar to the one described in the *Windows NT* section.

8.7. MacOS Box

8.7.1. MacOS 8/9

First of all, you need to open the TCP/IP Control Panel as shown below in the Apple menu.




Figure 8-12. Accessing The TCP/IP Control Panel

8.7.1.1. With an Automatic DHCP Configuration

If you configured your firewall to be a DHCP server, follow this very procedure, otherwise go to the next section.




Figure 8-13. Automatic Configuration of Internet Access For MacOS

In the dialog that appears, fill the fields as shown hereafter:

- Connect via: Ethernet;
- Configure: Using DHCP server;
- DHCP Client ID: 192.168.0.1.

8.7.1.2. For a Manual Configuration

If you have no DHCP server on your local network, follow this procedure:




Figure 8-14. Manual Configuration of Internet Access For MacOS

In the dialog that appears fill the fields as shown here:

- Connect via: Ethernet ;
- Configure: Manually ;
- IP address: 192.168.0.248 ;
- Subnet Mask: 255.255.255.0 ;
- Router Address: 192.168.0.1 ;
- Name Servers Addresses: 192.168.0.10 ; 192.168.0.11
- Search Domain: myco.com ;



The name server's addresses may be the addresses of the internal DNSs or those of your Internet Service Provider's servers.

8.7.2. MacTCP

1. In the MacTCP control panel, select the Ethernet network driver (caution, it's not EtherTalk) then click the More... button.
2. Under Gateway Address, enter the address of the *Linux* box sharing the connection (192.168.0.1 in our example).
3. Click OK to save the settings. You may have to restart your system to test these settings.

8.8. OS/2 Warp Box

The TCP/IP protocol should already be installed. If not, install it.


1. Go in Programs, then TCP/IP (LAN), then TCP/IP Settings.
2. Under Routing, choose Add. In Type, select default.
3. Fill the Router address field with the address of your *Linux* box sharing the Internet connection (192.168.0.1 in our example).
4. Now close the TCP/IP control panel, answer Yes to all questions, then reboot your system before testing the settings.

Chapter 9. Monitoring the Firewall


We will take a look at the monitoring tools available for your firewall system: System and Network usage graphics, and the most powerful tool when it comes to system auditing: the system logs.

9.1. System and Network Usage

The main screen of the Monitoring group shows the firewall's uptime (that is, for how long has the system been running without being rebooted); how many users are currently connected, and the system load averages for the past 1, 5 and 15 minutes.



9.1.1. System Usage Monitoring



This section's two graphics inform you about your system's load. They are good indicators of how well your system is performing with its actual usage, and can be used to support decisions regarding CPU/RAM upgrades.


- avgload: represents the CPU average usage for the last 24 hours. The unit used roughly indicates the number of processes trying to access the CPU at the same time. A normal load should remain below 2. If the load is between 2 and 5, your system is a rather busy one. Above 6, you should consider upgrading your CPU.
- memusage: represents your main RAM memory's usage (in megabytes). Different colors are used to give more precise information about the way memory is used (RAM used in black, RAM free in green, Swap in red, and cache in yellow).



By default, a daily graphic is shown. Clicking on  at the right of the graphic will show you the daily, weekly, monthly and yearly graphics in a single page. This can be handy to plan for system usage. Click on Back to return to the daily graphic.

Clicking on the Refresh button will update the graphs.

9.1.1.1. CPU Load Monitoring



This section holds graphics displaying CPU usage with different time scales.


CPU load average is shown graphically by Day, Week, Month and Year all in one page. The unit used roughly indicates the number of processes trying to access the CPU at the same time. Normal values are below 2. Values above 6 indicate that you should consider upgrading your CPU.

Click on the Refresh button to update the graphs.

Click on to return to the System Usage section.

Clicking on will take you to the default (starting) page.

9.1.1.2. Memory Usage Monitoring



This section holds graphics displaying RAM memory usage with different time scales.


RAM (physical) memory usage is shown graphically by Day, Week, Month and Year all in one page. Different colors are used to give more precise information about the way memory is used (RAM used in black, RAM free in green, Swap in red, and cache in yellow).

Clicking on the Refresh button will update the graphs.

Click on to return to the System Usage section.

Clicking on will take you to the default (starting) page.

9.1.2. Network Traffic Monitoring



The graphics shown here inform you about the inbound/outbound network traffic on your interfaces.

The first page shows the traffic for all interfaces during the last hour (by default). The units used will be adjusted according to the traffic on each interface, so you can have the traffic expressed in bytes/sec., kbytes/sec., mbytes/sec., etc.

At the top of the page is a list of the available time scales for the graphics: Hourly, Daily, Weekly, Monthly and Yearly. To change the time scale, just click on the corresponding link.

Each graphic also tells you the average and maximum inbound/outbound traffic for the network interfaces. Inbound traffic is represented in green and outbound traffic in dark gray.



Clicking on the magnifying glass icon at the right of each graphic (available only in Hourly Graph mode) will take you to traffic statistics on the corresponding interface.


Clicking on the Refresh button will update the graphics.

You could use the graphs shown above to plan your network connection times and bandwidth, for example.

9.2. Logs

System logs are very powerful tools for system auditing and performance analysis. Of course, they must be read by someone (the system administrator) to be useful. All events happening on your firewall system are logged; you can access logs under the Logs subsection.

9.2.1. System Messages Log



The System item allows you to take a look at the system log.

The table shows all system messages logged by the facility, for example: failed/accepted passwords for sshd, NTP updates, system maintenance scripts run, etc.

The table columns represent, respectively: at what date and time, which process, and which message was generated for each event logged.

Click on Refresh to get the latest entries.

9.2.2. Authentication Logs

	Time	Process	Message
May 9	11:01:16	msec	changed mode of /var/log/httpd-naat/ssl_cache.sem from 600 to 640
May 9	11:01:16	msec	changed mode of /etc/rc.d/init.d/squid from 711 to 700
May 9	11:01:16	msec	changed mode of /etc/rc.d/init.d/adsl from 644 to 700
May 9	11:01:16	msec	changed mode of /var/log/httpd/ssl_cache.sem from 600 to 640
May 9	11:01:16	msec	changed mode of /etc/rc.d/init.d/named from 711 to 700
May 9	11:01:16	msec	changed mode of /var/log/sudo.log from 600 to 640
May 9	11:02:23	sshd	Could not reverse map address 192.168.0.249.
May 9	11:02:27	sshd	Accepted password for peter from 192.168.0.249 port 35789 ssh2
May 9	12:01:14	msec	changed mode of /etc/rc.d/init.d/shorewall from 711 to 700
May 9	12:01:14	msec	changed mode of /etc/rc.d/init.d/prelude from 711 to 700
May 9	12:01:14	msec	changed mode of /etc/rc.d/init.d/squid from 711 to 700
May 9	12:01:14	msec	changed mode of /etc/rc.d/init.d/named from 711 to 700
May 9	13:01:09	msec	changed mode of /var/log/httpd-naat from 750 to 755
May 9	13:01:09	msec	changed mode of /etc/rc.d/init.d/shorewall from 711 to 700


The Authentication item allows you to take a look at your system's authentication security-related log.

The table describes the logged events related to authentication, such as file mode changes to log files, services starting/stopping, failed log-in attempts on the console and using sshd, etc. Typical services for this kind of messages are:

- msec: Mandrake's security tool;
- sshd: Secure shell daemon;
- xinetd: Secure replacement for inetd, the Internet daemon;

Click on Refresh to get the latest entries.

9.2.3. Firewall Logs



The Firewall item allows you to take a look at your system's packet filtering logs.

All firewall chain reports will be found in this section. Reports can be generated according to different criteria:

- Everything and name resolution: Shows all details about the packets, namely: packet number, start, interval, protocol, source IP, host name and port, destination IP, host name and port and packet options.
- Destination IP: Shows only the following details about the packets: packet number, start, interval and destination IP.
- Source IP: Shows only the following details about the packets: packet number, start, interval and source IP.
- Source and destination IP: Shows only the following details about the packets: packet number, start, interval, source IP and destination IP.
- With destination port: Shows only the following details about the packets: packet number, start, interval, source IP, destination IP and port.
- With source port: Shows only the following details about the packets: packet number, start, interval, source IP and port and destination IP.
- With source and destination port: shows only the following details about the packets: packet number, start, interval, source IP and port and destination IP and port.
- With TCP options: Shows the same details as the "Everything and name resolution" criteria, except the source and destination host names.



Clicking on  at the left of each of the above items will show the corresponding Firewall Logs Summary window, for example:

```
Generated Mon Apr 15 11:16:09 ART 2002 by root.
5 of 456 items in the file "/var/log/messages" are packet logs, one has unique characteristics.
First packet log entry: Apr 15 10:53:26; last: Apr 15 10:53:26.
All entries where logged by the same host: "e500".
All entries are from the same chain: "Shorewall:fw2all:REJECT:".
All entries have the same target: "-".
All entries are from the same interface: "".
Only entries with a count larger than 2 are shown.
```


After the above messages follows a table with the packet details.



This logs might not be immediately available due to system activity.

Click on Refresh to get the latest entries.

9.2.4. Prelude IDS Logs



The Prelude Intrusion Detection System (IDS) item allows you to take a look at your system’s Prelude IDS security-related logs.

The Prelude IDS reports abnormal (non-expected, suspicious, bizarre, etc.) packets targeted to your network. It also tries to avoid receiving such attacks.

If the Prelude IDS is not active on your system, the report will show something like:

```
Reports
empty Prelude IDS was not activated yet
```

If the Prelude IDS is active but no attacks have been attempted, on your system the report will show something like:

```
Reports
empty No log available!
```




Some logs might not be immediately available due to system activity.



When logs become available, clicking on will show the Prelude IDS Logs Summary window. Click on Refresh to get the latest entries.

9.2.5. Snort IDS Logs



The Snort Intrusion Detection System (IDS) item allows you to take a look at your system’s Snort IDS security-related logs.

The Snort IDS analyzes traffic coming into your network looking for matches against pre-defined rules, and performs actions accordingly.

If the Snort IDS is not active on your system, the report will show something like:

Reports

empty	Snort IDS was not activated yet
-------	---------------------------------

If the Snort IDS is active but no attacks have been attempted on your system, the report will show something like:

Reports

empty	No log available!
-------	-------------------



Some logs might not be immediately available due to system activity.



When logs become available, clicking on will show the Snort IDS Logs Summary window.

Click on Refresh to get the latest entries.

Below, you will find a sample *Snort* report.





Figure 9-1. Sample Snort Report

9.2.6. Proxy Server Logs



The Web Proxy item allows you to take a look at your system's Squid proxy server logs.

Squid is a high performance proxy caching server for web clients which supports FTP, gopher and HTTP data objects. It also caches DNS lookups. These make for more effective use of your Internet bandwidth together with more responsive web clients.

This page holds access information, resources (memory, disk space) consumption and configuration errors for the Squid web proxy server.

If the Squid proxy server is not active on your system, the report will show something like:

Reports

empty The WebProxy server was not activated yet

If the proxy server is active, but no events were logged, the report will show something like:

Empty list...




Some logs might not be immediately available due to system activity.



When logs become available, clicking on  will show the Proxy Server Logs Summary window. Click on Refresh to get the latest entries.

9.2.7. DHCP Logs



The DHCP item allows you to take a look at your system’s DHCP server logs.

DHCP server messages, like IP assignment to interfaces, DHCP packets from clients, and the like, are shown here.

If the DHCP server is not active on your system, the report shows something like:

Reports

empty The DHCP server was not activated yet



If the DHCP server is active on your system, then clicking on  will show you a table with the DHCP Subnet Information. The table has the following columns: Location, Subnet, Netmask, IP range, Router, IPs defined, IPs used and IPs free.

Click on Refresh to get the latest entries.

Chapter 10. Management Tools

We will take a look at the management tools available for your firewall system: remote console, backup/restore of the firewall configuration, and software updates.





Figure 10-1. Tools' Main Screen

When you click on the Tools link, you are presented with the screen shown in figure 10-1 where you will see a Shutdown the Firewall button. If you click on that button, the firewall system will be shut down, thus leaving all users who depend on it without connection to the Internet, for example. Use this button with care.

10.1. Remote Connection Using SSH



After entering the IP address (or hostname) of the host you want to connect to, a window running an SSH console will appear. It will let you perform some actions as if you were sitting at that system's console.

Usually, the system you want to connect to is the firewall itself. At the prompt, enter "admin" as the login (without the quotes) and then, admin's password:

```
firewall login:admin
```

```
admin@firewall's password: *****
```


Upon successful connection, you will be faced with the shell and will be able to perform administration tasks as if you were sitting directly in front of the firewall machine.



Take into account that the actions performed using the remote console may cause your modifications to be lost. If later you decide to use again the web interface to modify the configuration, the manual modifications made with the console will be lost. Thus, make sure to use the web interface whenever possible, rather than the SSH connection, to change your firewall's parameters.

Finally, note that if a SSH client is installed on your machine, you can connect to the firewall host directly and bypass the web interface. Please, bear in mind that the previous warning is still valid.

10.2. Backup and Restore



This feature backs up the entire configuration of your firewall system, letting you quickly recover in case of a major system failure. It will also help you to easily reconfigure a new firewall system based on the backed-up configuration.

To create a backup of your configuration file, click on the Backup button. This will build the backup and take you to a page where you can retrieve it. Please refer to that page’s help for instructions on retrieval.

To restore the configuration, you need to choose the appropriate configuration file by clicking on the "Browse..." button. A window will pop up, allowing you to select the file containing the backed-up configuration. If you followed the procedure described in the backup help page, this should be "/mnt/floppy/ConfigurationBackup" (without the quotes). If you do not save your backups to floppy, then you are advised to copy them to floppy regularly and store the floppies in a safe place.



This backup will only contain the parameters handled by the web interface. If you have made modifications to some configuration parameters by other means, e.g. the Remote Login tool or a direct connection through SSH, those specific modifications (files) will have to be backed up manually.

Choose the Configuration file	/mnt/floppy/ConfigurationBackup
-------------------------------	---------------------------------

When done, click on the Upload button. You will be taken to a confirmation page to apply/modify/cancel the changes. Please refer to that page’s help for more details.

Below is a sample screen of the Restore section completed with the file name of the backed up configuration file to restore on your firewall system.




Figure 10-2. Sample Restore Screen

Once you have clicked on the Upload button, you will get a confirmation screen, like the one shown in figure 10-3 click on the Apply button to apply the configuration to your firewall system.





Figure 10-3. Apply Configuration From The Restored File

10.2.1. Retrieving the Backed-up Configuration



Through this page, you can store your firewall configuration onto a floppy disk (or another removable medium).


After having clicked on the Backup button in the previous page you are provided with a link (named "Backup file") in order to be able to retrieve the backed-up configuration. Please, proceed as follows (we assume that you want to store your configuration onto a floppy disk):

- Insert a blank floppy disk onto your floppy disk drive.
- Right-click (depending on your web browser, you might need to shift-click) on the "Backup file" link and select the "Download link" (again, this might be different for your web browser) option.
- When prompted for the file name, enter `"/mnt/floppy/ConfigurationBackup"` (without the quotes). Of course, you can change "ConfigurationBackup" for another file name.
- Accept and save it.

After performing the above steps, your firewall configuration will be on the floppy disk. Remove it from the floppy drive, write-protect it and store it in a safe place. Then click on the **Next** button to return to the Backup/Restore page.

Clicking on **Cancel** will take you to the default (starting) page.

10.3. Update Software



This wizard will let you perform updates of all packages installed on your system. For security reasons, it is essential that you check regularly for software updates.

Current Update Mirror	ftp://ftp-linux.cc.gatech.edu/pub/linux/distributions/mandrake/updates
------------------------------	--

This shows the currently selected mirror site for retrieving software updates.

Modify Update Mirror type	Registered Mirror
----------------------------------	-------------------

You have to select between three different types of mirrors:

- **Registered Mirror:** this option will give you access to the updates site dedicated to MandrakeSecurity. It contains updates for the regular base Linux distribution upon which resides MandrakeSecurity, but also MandrakeSecurity specific packages, as well as optional modules.
- **Official Mirror:** on the following page, you are given a list of all Mandrake Linux official mirror sites. It is strongly recommended that you select one of these as your mirror.
- **Personal Mirror:** on the following page, you will be able to manually enter the URL of the site hosting the updates for your system.

Make your selection and click on **Next** to continue to the next step. Please refer to the next step's help page for more details.

Clicking on **Cancel** will take you to the default (starting) page cancelling this wizard.

Security And Network Issues

Up until now in this manual, you have been reading very practical information. You should be able to efficiently configure your server and be happy with it by now.

However, all that is just a glimpse of your **Mandrake Linux** system's possibilities. In order to dig in a little deeper, we chose to add two chapters to complete your **Mandrake Linux** knowledge:

- *"Security Under GNU/Linux"*, page 135: this is a must read for any system administrator. Even though you can make your **Mandrake Linux** system quite secure with default tools, efficient security can only be achieved through active administration, taking care of both physical and logical global system security;
- *"Networking Overview"*, page 173: a server is meant to bring services to a network. This manual would have been incomplete without a chapter dedicated to networking. The configuration of the network itself and the different protocols are tackled.

Chapter 11. Security Under GNU/Linux

This document is a general overview of security issues that face the administrator of *GNU/Linux* systems. It covers general security philosophy and a number of specific examples of how to better secure your *GNU/Linux* system from intruders. Also included are pointers to security-related material and programs.



1. The original document (see below) has been adapted to **Mandrake Linux** distribution, removing parts, changing others, etc.

11.1. Preamble

This chapter is based on a *HOWTO* by Kevin Fenzi and Dave Wreski which original is hosted by the Linux Documentation Project (<http://www.tldp.org>).

11.1.1. Copyright Information

This document is copyrighted (c) 1998 - 2002 Kevin Fenzi and Dave Wreski.

Modifications from v1.3.1, 11 February 2002, (C)opyright 2000-2002 MandrakeSoft and distributed under the following terms:

- Linux HOWTO documents may be reproduced and distributed in whole or in part, in any medium, physical or electronic, as long as this copyright notice is retained on all copies. Commercial redistribution is allowed and encouraged; however, the authors would like to be notified of any such distributions.
- All translations, derivative works, or aggregate works incorporating any Linux HOWTO documents must be covered under this copyright notice. That is, you may not produce a derivative work from a HOWTO and impose additional restrictions on its distribution. Exceptions to these rules may be granted under certain conditions; please contact the Linux HOWTO coordinator at the address given below.
- If you have questions, please contact Tim Bynum, the Linux HOWTO coordinator, at tjbynum@metalab.unc.edu.

11.1.2. Introduction

This chapter covers some of the main issues that affect *GNU/Linux* security. General philosophy and net-born resources are discussed.

A number of other *HOWTO* documents overlap with security issues, and those documents have been pointed to wherever appropriate.

This chapter is **not** meant to be an up-to-date exploits document. Large numbers of new exploits happen all the time. This chapter will tell you where to look for such up-to-date information, and will give you some general methods to prevent such exploits from taking place.

11.2. Overview

This chapter will attempt to explain some procedures and commonly-used software to help your *GNU/Linux* system be more secure. It is important to discuss some of the basic concepts first, and create a security foundation, before we get started.

11.2.1. Why Do we Need Security?

In the ever-changing world of global data communications, inexpensive Internet connections, and fast-paced software development, security is becoming more and more of an issue. Security is now a basic requirement because global computing is inherently insecure. As your data goes from point A to point B on the Internet, for example, it may pass through several other points along the way, giving other users the opportunity to intercept, or worst, alter it. Even other users on your system may maliciously transform your data into something you did not intend for. Unauthorized access to your system may be obtained by intruders, also known as “crackers”, who then use advanced knowledge to impersonate you, steal information from you, or even deny you access to your own resources. If you’re wondering what the difference is between a “hacker” and a “cracker”, see Eric Raymond’s document, How to Become a Hacker (<http://www.tuxedo.org/~esr/faqs/hacker-howto.html>).

11.2.2. How Secure Is Secure?

First, keep in mind that no computer system can ever be completely secure. All you can do is make it increasingly difficult for someone to compromise your system. For the average home *GNU/Linux*, not much is required to keep the casual cracker at bay. However, for high profile *GNU/Linux* users (banks, telecommunications companies, etc.), much more work is required.

Another factor to take into account is that the more secure your system is, the more intrusive your security becomes. You need to decide where in this balancing act your system will still be usable, and yet secure for your purposes. For instance, you could require everyone dialing into your system to use a call-back modem to call them back at their home number. This is more secure, but if someone is not at home, it makes it difficult for them to log in. You could also set up your *GNU/Linux* system with no network nor connection to the Internet, but this limits its usefulness.

If you are a medium-to-large-size site, you should establish a security policy stating how much security is required by your site and what auditing is in place to check it. You can find a well-known security policy example at faqs.org (<http://www.faqs.org/rfcs/rfc2196.html>). It has been recently updated, and contains a great framework for establishing a security policy for your company.

11.2.3. What Are You Trying to Protect?

Before you attempt to secure your system, you should determine what level of threat you have to protect against, what risks you should or should not take, and how vulnerable your system is as a result. You should analyze your system to know what you’re protecting, why you’re protecting it, what value it has, and who has responsibility for your data and other assets.

- **Risk** is the possibility that an intruder may be successful in attempting to access your computer. Can an intruder read or write files, or execute programs that could cause damage? Can they delete critical data? Can they prevent you or your company from getting important work done? Don’t forget: someone gaining access to your account, or your system, can also impersonate you.

Additionally, having one insecure account on your system can result in your entire network being compromised. If you allow a single user to log in using a `.rhosts` file, or to use an insecure service, such as `tftp`, you risk an intruder getting “his foot in the door”. Once the intruder has a user account on your system, or someone else’s system, it can be used to gain access to another system, or another account.

- **Threat** is typically from someone with motivation to gain unauthorized access to your network or computer. You must decide who you trust to have access to your system, and what threat they could pose.

There are several types of intruders, and it is useful to keep their different characteristics in mind as you are securing your systems.

- **The Curious** – This type of intruder is basically interested in finding out what type of system and data you have.
- **The Malicious** – This type of intruder is out to either bring down your systems, or deface your web page, or otherwise force you to spend time and money recovering from the damage he has caused.

- **The High-Profile Intruder** – This type of intruder is trying to use your system to gain popularity and infamy. He might use your high-profile system to advertise his abilities.
 - **The Competition** – This type of intruder is interested in what data you have on your system. It might be someone who thinks you have something that could benefit him, financially or otherwise.
 - **The Borrowers** – This type of intruder is interested in setting up shop on your system and using its resources for their own purposes. He typically will run chat or IRC servers, porn archive sites, or even DNS servers.
 - **The Leapfrogger** – This type of intruder is only interested in your system to use it to get into other systems. If your system is well connected or a gateway to a number of internal hosts, you may well see this type trying to compromise your system.
- Vulnerability describes how well-protected your computer is from another network, and the potential for someone to gain unauthorized access.

What's at stake if someone breaks into your system? Of course the concerns of a dynamic PPP home user will be different from those of a company connecting their machine to the Internet, or another large network.

How much time would it take to retrieve/recreate any data that was lost? An initial time investment now can save ten times more time later if you have to recreate data that was lost. Have you checked your backup strategy and verified your data lately?

11.2.4. Developing a Security Policy

Create a simple, generic policy for your system that your users can readily understand and follow. It should protect the data you're safeguarding as well as the privacy of the users. Some things to consider adding are: who has access to the system (can my friend use my account?), who's allowed to install software on the system, who owns what data, disaster recovery, and appropriate use of the system.

A generally-accepted security policy starts with the phrase:

“ That which is not permitted is prohibited. ”

This means that unless you grant access to a service for a user, that user shouldn't be using that service until you do grant access. Make sure the policies work on your regular user account. Saying, “Ah, I can't figure out this permissions problem, I'll just do it as root”, can lead to security holes that are very obvious, and even ones that haven't been exploited yet.

RFC 1244 (<http://www.faqs.org/rfcs/rfc1244.html>) is a document that describes how to create your own network security policy.

rfc1281 (<http://www.faqs.org/rfcs/rfc1281.html>) is a document that shows a security policy example with detailed descriptions of each step.

Finally, you might want to look at the COAST policy archive (<ftp://coast.cs.purdue.edu/pub/doc/policy>) to see what some real-life security policies look like.

11.2.5. Means of Securing your Site

This section will discuss various means with which you can secure the assets you have worked hard for: your local computer, your data, your users, your network, even your reputation. What would happen to your reputation if an intruder deleted some of your users' data? Or defaced your web site? Or published your company's corporate project plan for the next quarter? If you are planning a network installation, there are many factors you must take into account before adding a single computer to your network.

Even if you have a single dial up PPP account, or just a small site, this does not mean intruders won't be interested in your systems. Large, high-profile sites are not the only targets – many intruders simply want to exploit as many sites as possible, regardless of their size. Additionally, they may use a security hole in your site to gain access to other sites you're connected to.

Intruders have a lot of time on their hands, and can avoid guessing how you've obscured your system just by trying all the possibilities. There are also a number of reasons an intruder may be interested in your systems, which we will discuss later.

11.2.5.1. Host Security

Perhaps the area of security on which administrators concentrate most is host-based security. This typically involves making sure your own system is secure, and hoping everyone else on your network does the same. Choosing good passwords, securing your host's local network services, keeping good accounting records, and upgrading programs with known security exploits are among the things the local security administrator is responsible for doing. Although this is absolutely necessary, it can become a daunting task once your network becomes larger than a few computers.

11.2.5.2. Local Network Security

Network security is as necessary as local host security. With hundreds, thousands, or more computers on the same network, you can't rely on each one of those systems being secure. Ensuring that only authorized users can use your network, building firewalls, using strong encryption, and ensuring there are no "rogue" (that is, unsecured) computers on your network are all part of the network security administrator's duties.

This document will discuss some of the techniques used to secure your site, and hopefully show you some of the ways to prevent an intruder from gaining access to what you are trying to protect.

11.2.5.3. Security Through Obscurity

One type of security that must be discussed is "security through obscurity". This means, for example, moving a service that has known security vulnerabilities to a non-standard port in hopes that attackers won't notice it's there and thus won't exploit it. Rest assured that they can determine that it's there and will exploit it. Security through obscurity is no security at all. Simply because you may have a small site, or a relatively low profile, does not mean an intruder won't be interested in what you have. We'll discuss what you're protecting in the next sections.

11.2.6. Organization of This Chapter

This chapter has been divided into a number of sections. They cover several broad security issues. The first, *Physical Security*, page 138, covers how you need to protect your physical machine from tampering. The second, *Local Security*, page 143, describes how to protect your system from tampering by local users. The third, *Files and File-System Security*, page 144, shows you how to set up your file systems and permissions on your files. The next, *Password Security and Encryption*, page 149, discusses how to use encryption to better secure your machine and network. *Kernel Security*, page 154 discusses what kernel options you should set or be aware of for a more secure system. *Network Security*, page 157, describes how to better secure your *GNU/Linux* system from network attacks. *Security Preparation (Before You Go On-Line)*, page 164, discusses how to prepare your machine(s) before bringing them on-line. Next, *What to Do During and After a Breach*, page 165, discusses what to do when you detect a system compromise in progress or detect one that has recently happened. In *Security Sources*, page 167, some primary security resources are enumerated. The Q & A section *Frequently Asked Questions*, page 169, answers some frequently-asked questions, and finally, a conclusion in *Conclusion*, page 170.

The two main points to realize when reading this chapter are:

- Be aware of your system. Check system logs such as `/var/log/messages` and keep an eye on your system;
- Keep your system up-to-date by making sure you have installed the current versions of software and have upgraded per security alerts. Just doing this will help make your system markedly more secure.

11.3. Physical Security

The first layer of security you need to take into account is the physical security of your computer systems. Who has direct physical access to your computer? Should they? Can you protect your computer from their tampering? Should you?

How much physical security you need on your system is very dependent on your situation, and/or budget.

If you are a home user, you probably don't need a lot (although you might need to protect your computer from tampering by children or annoying relatives). If you are in a lab, you need considerably more, but users will still need to be able to get work done on the computers. Many of the following sections will help out. If you are in an office, you may or may not need to secure your computer off-hours or while you are away. At some companies, leaving your console unsecured is a termination offense.

Obvious physical security methods such as locks on doors, cables, locked cabinets, and video surveillance are all good ideas, but beyond the scope of this chapter :-)

11.3.1. Computer Locks

Many modern PC cases include a "locking" feature. Usually this will be a socket on the front of the case that allows you to turn an included key to a locked or unlocked position. Case locks can help prevent someone from stealing your PC, or opening up the case and directly manipulating/stealing your hardware. They can also sometimes prevent someone from rebooting your computer from their own floppy or other hardware.

These case locks do different things according to the support in the motherboard and how the case is constructed. On many PC's, they make it so you have to break the case to get the case open. On some others, they will not let you plug in new keyboards or mice. Check your motherboard or case instructions for more information. This can sometimes be a very useful feature, even though the locks are usually very low quality and can easily be defeated by attackers with locksmithing.

Some computers (most notably SPARCs and Macs) have a dangle on the back: if you put a cable through attackers would have to cut the cable or break the case to get into it. Just putting a padlock or combo lock through these can be a good deterrent to someone stealing your computer.

11.3.2. BIOS Security

The *BIOS* is the lowest level of software that configures or manipulates your x86-based hardware. *grub* and other *GNU/Linux* boot methods access the *BIOS* to determine how to boot up your *GNU/Linux* computer. Other hardware that *GNU/Linux* runs on has similar software (Open Firmware on Macs and new Suns, Sun boot PROM, etc...). You can use your *BIOS* to prevent attackers from rebooting your computer and manipulating your *GNU/Linux* system.

Many PC *BIOS*s let you set a boot password. This doesn't provide all that much security (the *BIOS* can be reset, or removed if someone can get into the case), but might be a good deterrent (i.e. it will take time and leave traces of tampering). Similarly, on *S/Linux* (*GNU/Linux* for *SPARC*^(™) processor computers), your EEPROM can be set to require a boot-up password. This might slow attackers down.

Another risk of trusting *BIOS* passwords to secure your system is the default password problem. Most *BIOS* makers don't expect people to open up their computer and disconnect batteries if they forget their password and have equipped their *BIOS*es with default passwords that work regardless of your chosen password. Some of the more common passwords include:

```
j262 AWARD_SW AWARD_PW lkwpete Biostar AMI Award
  bios BIOS setup cmos AMI!SW1 AMI?SW1 password hewittrand shift + s y
  x z
```

I tested an Award *BIOS* and AWARD_PW worked. These passwords are quite easily available from manufacturers' web sites and astalavista (<http://astalavista.box.sk>) and as such a *BIOS* password cannot be considered adequate protection from a knowledgeable attacker.

Many *x86 BIOS*s also allow you to specify various other good security settings. Check your *BIOS* manual or look at it the next time you boot up. For example, some *BIOS*s disallow booting from floppy drives and some require passwords to access some *BIOS* features.



If you have a server computer, and you set up a boot password, your computer will not boot up unattended. Keep in mind that you will need to come in and supply the password in the event of a power failure ;(

11.3.3. OpenBoot Security

The *PROM* is the lowest level of software that configures or manipulates your SPARC-based hardware. *SIL0* and other *GNU/Linux* boot methods access the *PROM* to determine how to boot up your *GNU/Linux* computer. Other hardware that *GNU/Linux* runs on has similar software (OpenFirmware on Macs and new Suns, *x86 BIOS*, etc.). You can use your *PROM* to prevent attackers from rebooting your computer and manipulating your *GNU/Linux* system.

OpenBoot is much more advanced than a PC *BIOS* when it comes to security (consult the “*Installation Guide*” on how to access and use *OpenBoot*).

1.



It is important to set your password before setting the security mode, as you would be unable to set it any more. Moreover, SUN claims you need to contact your vendor's customer support service to make your computer bootable again.

This is an interaction example on how to set your boot password:

```
> password
> New password (only first 8 chars are used):
> Retype new password:
>
```

2. You can choose between three security levels setting the `security-mode` variable:

- a. Full: all commands except for `go` require the password.
- b. Command: all commands except for `boot` and `go` require the password.
- c. None: no password required (default).

This is an interaction example on how to set your security mode:

```
> setenv security-mode full
>
```



If you have a server computer, and you set up a boot password, your computer will not boot up unattended. Keep in mind that you will need to come in and supply the password in the event of a power failure ;(

11.3.4. Boot Loader Security

Keep in mind when setting all these passwords that you need to remember them :-). Also remember that these passwords will only slow the determined attacker. They won't prevent someone from booting from a floppy and mounting your root partition.

If you are using security in conjunction with a boot loader, you might as well disable booting from a floppy in your computer's *BIOS*, and password-protect the *BIOS*.

Also keep in mind that the `/etc/lilo.conf` will need to be mode 600 (readable and writing for root only), or others will be able to read your boot passwords!

If you are using security in conjunction with a boot loader, you might as well password-protect the *PROM*.



Once again, if you have a server computer, and you set up a boot password, your computer will not boot up unattended. Keep in mind that you will need to come in and supply the password in the event of a power failure!

11.3.4.1. With GRUB

The various *GNU/Linux* boot loaders also can have a boot password set. *grub* is quite flexible in that sense: your default configuration file `/boot/grub/menu.lst` may contain a line allowing the loading of a new configuration file with different options (this new file may contain a new password to access another third configuration file and so on).

So you must add a line in your `/boot/grub/menu.lst` file, something like:

```
password very_secret /boot/grub/menu2.lst
```

and of course generate a new `/boot/grub/menu2.lst` configuration file where you move insecure entries previously removed from `/boot/grub/menu.lst`.

>From the grub info page:

- Command: `password passwd new-config-file`
 Disable all interactive editing control (menu entry editor and command line). If the password `PASSWD` is entered, it loads the `NEW-CONFIG-FILE` as a new config file and restarts the GRUB Stage 2.

11.3.4.2. With LILO

LILO has password and restricted settings; password requires password at boot time, whereas restricted requires a boot-time password only if you specify options (such as `single`) at the *LILO* prompt.

From the `lilo.conf` man page:

```
password=password
    The per-image option 'password=...' (see below)
    applies to all images.

restricted
    The per-image option 'restricted' (see below)
    applies to all images.

password=password
    Protect the image by a password.

restricted
    A password is only required to boot the image if
    parameters are specified on the command line
    (e.g. single).
```

11.3.4.3. With SILO

The *SILO* boot loader may also have a boot password: `password` requires password at boot time, whereas `restricted` requires a boot-time password only if you specify options (such as `single`) at the *SILO* prompt.

From the `sil0.conf` man page:

```
password=password
    Protect booting by a password. The password is
    given in cleartext in the configuration file.
    Because of that the configuration file should be
    only readable by the super user and the password
    should differ if possible from other passwords on
    the system.

restricted
    A password is only required to boot the image spec-
    ified in /etc/silo.conf if parameters are specified
    on the command line or if the image is not speci-
    fied in the configuration file at all (i.e. arbi-
    trary file load).
```

11.3.5. `xlock` and `vlock`

If you wander away from your computer from time to time, it is nice to be able to “lock” your console so that no one can tamper with or look at your work. Two programs that do this are: `xlock` and `vlock`.

`xlock` is a *X* display locker. You can run `xlock` from any *xterm* on your console and it will lock the display and require your password to unlock. Most desktop environment also propose this feature in their respective menus.

`vlock` is a simple little program that allows you to lock some or all of the virtual consoles on your *GNU/Linux* box. You can lock just the one you are working in or all of them. If you just lock one, others can come in and use the console; they will just not be able to use your virtual console until you unlock it.

Of course, locking your console will prevent someone from tampering with your work, but won’t prevent them from rebooting your computer or otherwise disrupting your work. It also does not prevent them from accessing your computer from another computer on the network and causing problems.

More importantly, it does not prevent someone from switching out of the *X Window System* entirely, and going to a normal virtual console login prompt, or to the VC that *X11* was started from, and suspending it, thus obtaining your privileges. For this reason, you might consider only using it while under control of *KDM* (or other).

11.3.6. Security of Local Devices

If you have a webcam or a microphone attached to your system, you should consider if there is some danger of an attacker gaining access to those devices. When not in use, unplugging or removing such devices might be an option. Otherwise you should carefully read and look at any software with provides access to such devices.

11.3.7. Detecting Physical Security Compromises

The first thing to always note is when your computer was rebooted. Since *GNU/Linux* is a robust and stable OS, the only times your computer should reboot is when **you** take it down for OS upgrades, hardware swapping, or the like. If your computer has rebooted without you doing it, that may be a sign that an intruder has compromised it. Many of the ways that your computer can be compromised require the intruder to reboot or power off your computer.

Check for signs of tampering on the case and computer area. Although many intruders clean traces of their presence out of logs, it’s a good idea to check through them all and note any discrepancy.

It is also a good idea to store log data at a secure location, such as a dedicated log server within your well-protected network. Once a computer has been compromised, log data becomes of little use as it most likely has also been modified by the intruder.

The *syslog* daemon can be configured to automatically send log data to a central *syslog* server, but this is typically sent in unencrypted, allowing an intruder to view data as it is being transferred. This may reveal information about your network that is not intended to be public. There are *syslog* daemons available that encrypt the data as it is being sent.

Also be aware that faking *syslog* messages is easy – with an exploit program having been published. *syslog* even accepts net log entries claiming to come from the local host without indicating their true origin.

Some things to check for in your logs:

- short or incomplete logs;
- logs containing strange timestamps;
- logs with incorrect permissions or ownership;
- records of reboots or restarting of services;
- missing logs;
- su entries or logins from strange places.

We will discuss system log data *Keep Track of your System Accounting Data*, page 165 in this chapter.

11.4. Local Security

The next thing to take a look at is the security in your system against attacks from local users. Did we just say **local** users? Yes!

Getting access to a local user account is one of the first things that system intruders attempt while on their way to exploiting the root account. With lax local security, they can then “upgrade” their normal user access to root access using a variety of bugs and poorly set up local services. If you make sure your local security is tight, then the intruder will have another hurdle to jump.

Local users can also cause a lot of havoc with your system even (especially) if they really are who they say they are. Providing accounts to people you don’t know or for whom you have no contact information for is a very bad idea.

11.4.1. Creating New Accounts

You should make sure you provide user accounts with only the minimal requirements for the task they need to do. If you provide your son (age 10) with an account, you might want him to only have access to a word processor or drawing program, but be unable to delete data that is not his.

Several good rules of thumb when allowing other people legitimate access to your *GNU/Linux* computer:

- give them the minimal amount of privileges they need;
- be aware when/where they log in from, or should be logging in from;
- make sure you remove inactive accounts, which you can determine by using the `last` command and/or checking log files for any activity by the user;
- the use of the same `userid` on all computers and networks is advisable to ease account maintenance, and permits easier analysis of log data;
- the creation of group `user-ids` should be absolutely prohibited. User accounts also provide accountability, and this is not possible with group accounts.

Many local user accounts that are utilized in security compromises have not been used in months or years. Since no one is using them, they provide the ideal attack vehicle.

11.4.2. Root Security

The most sought-after account on your computer is the root (superuser) account. It has authority over the entire computer, which may also include authority over other computers on the network. Remember that you should only use the root account for very short, specific tasks, and should mostly run as a normal user. Even small mistakes made while logged in as the root user can cause problems. The less time you are on with root privileges, the safer you will be.

Several tricks to avoid messing up your own box as root:

- When doing some complex command, try running it first in a non-destructive way... especially commands that use globbing: e.g., if you want to do `rm -f foo*.bak`, first do `ls foo*.bak` and make sure you are going to delete the files you think you are. Using `echo` in place of destructive commands also sometimes works.
- Only become root to do single specific tasks. If you find yourself trying to figure out how to do something, go back to a normal user *shell* until you are **sure** what needs to be done by root.
- The command path for the root user is very important. The command path (that is, the `PATH` environment variable) specifies the directories in which the *shell* searches for programs. Try to limit the command path for the root user as much as possible, and **never** include `.` (which means “the current directory”) in your `PATH`. Additionally, never have writable directories in your search path, as this can allow attackers to modify or place new binaries in your search path, allowing them to run as root the next time you run that command.
- Never use the `rlogin/rsh/rexec` suite of tools (called the “r-utilities”) as root. They are subject to many sorts of attacks, and are downright dangerous when run as root. Never create a `.rhosts` file for root.
- The `/etc/securetty` file contains a list of terminals that root can login from. By default, this is set to only the local virtual consoles (ttys). Be very wary of adding anything else to this file. You should be able to log in remotely as your regular user account and then `su` if you need to (hopefully over `ssh` or other encrypted channel), so there is no need to be able to login directly as root.
- Always be slow and deliberate running as root. Your actions could affect a lot of things. Think before you type!

If you absolutely, positively need to allow someone (hopefully very trusted) to have root access to your computer, there are a few tools that can help. `sudo` allows users to use their password to access a limited set of commands as root. This would allow you to, for instance, let a user be able to eject and mount removable media on your *GNU/Linux* box, but have no other root privileges. `sudo` also keeps a log of all successful and unsuccessful `sudo` attempts, allowing you to track down who used what command to do what. For this reason, `sudo` works well even in places where a number of people have root access, because it helps you to keep track of changes made.

Although `sudo` can be used to give specific users special privileges for particular tasks, it does have several shortcomings. It should be used only for a limited set of tasks, like restarting a server, or adding new users. Any program that offers a *shell* escape will give root access to a user invoking it via `sudo`. This includes most editors, for example. Also, a program as innocuous as `/bin/cat` can be used to overwrite files, which could allow root to be exploited. Consider `sudo` as a means for accountability, and don’t expect it to replace the root user and still be secure.

11.5. Files and File-System Security

A few minutes of preparation and planning ahead before putting your systems on-line can help protect them and the data stored in them.

- There should never be a reason for users’ home directories to allow SUID/SGID programs to be run from there. Use the `nosuid` option in `/etc/fstab` for partitions that are writable by others than root. You may also wish to use `nodev` and `noexec` on users’ home partitions, as well as `/var`, thus prohibiting execution of programs, and creation of character or block devices, which should never be necessary anyway.

- If you are exporting file systems using NFS, be sure to configure `/etc/exports` with the most restrictive access possible. This means not using wildcards, not allowing root write access, and exporting read-only wherever possible.
- Configure your users' file-creation `umask` to be as restrictive as possible. See *umask Settings*, page 146.
- If you are mounting file systems using a network file system such as NFS, be sure to configure `/etc/fstab` with suitable restrictions. Typically, using `nodev`, `nosuid`, and perhaps `noexec`, are desirable.
- Set file system limits instead of allowing unlimited as default. You can control the per-user limits using the resource-limits PAM module and `/etc/pam.d/limits.conf`. For example, limits for group users might look like this:

```
@users    hard   core    0
@users    hard   nproc   50
@users    hard   rss     5000
```

This says to prohibit the creation of core files, restrict the number of processes to 50, and restrict memory usage per user to 5MB.

You can also use the `/etc/login.defs` configuration file to set the same limits.

- The `/var/log/wtmp` and `/var/run/utmp` files contain the login records for all users on your system. Their integrity must be maintained because they can be used to determine when and from where a user (or potential intruder) has entered your system. These files should also have 644 permissions, without affecting normal system operation.
- The immutable bit can be used to prevent accidentally deleting or overwriting a file that must be protected. It also prevents someone from creating a hard link to the file. See the `chattr(1)` man page for information on the immutable bit.
- `suid` and `SGID` files on your system are a potential security risk, and should be monitored closely. Because these programs grant special privileges to the user who is executing them, it is necessary to ensure that insecure programs are not installed. A favorite trick of crackers is to exploit `SUID-root` programs, then leave a `SUID` program as a backdoor to get in the next time, even if the original hole is plugged.

Find all `SUID/SGID` programs on your system, and keep track of what they are, so you are aware of any changes which could indicate a potential intruder. Use the following command to find all `SUID/SGID` programs on your system:

```
root# find / -type f \( -perm -04000 -o -perm -02000 \)
```

You can remove the `suid` or `SGID` permissions on a suspicious program with `chmod`, then restore them back if you absolutely feel it is necessary.

- World-writable files, particularly system files, can be a security hole if a cracker gains access to your system and modifies them. Additionally, world-writable directories are dangerous, since they allow a cracker to add or delete files as he wishes. To locate all world-writable files on your system, use the following command:

```
root# find / -perm -2 ! -type l -ls
```

and be sure you know why those files are writable. In the normal course of operation, several files will be world-writable, including some from `/dev`, and symbolic links, thus the `! -type l` which excludes these from the previous `find` command.

- Un-owned files may also be an indication that an intruder has accessed your system. You can locate files on your system that have no owner or belong to no group with the command:

```
root# find / -nouser -o -nogroup -print
```

- Finding `.rhosts` files should be a part of your regular system administration duties, as they should not be permitted on your system. Remember, a cracker only needs one insecure account to potentially gain access to your entire network. You can locate all `.rhosts` files on your system with the following command:

```
root# find /home -name .rhosts -print
```

- Finally, before changing permissions on any system files, make sure you understand what you are doing. Never change permissions on a file because it seems like the easy way to get things working. Always determine why the file has that permission before changing it.

11.5.1. umask Settings

The `umask` command can be used to determine the default file-creation mode on your system. It is the octal complement of the desired file mode. If files are created without any regard to their permission settings, the user could inadvertently give read or write permission to someone that should not have it. Typical `umask` settings include 022, 027, and 077 (which is the most restrictive). Normally, the `umask` is set in `/etc/profile`, so it applies to all users on the system. The file creation mask can be calculated by subtracting the desired value from 777. In other words, a `umask` of 777 would cause newly-created files to contain no read, write or execute permission for anyone. A mask of 666 would cause newly-created files to have a mask of 111. For example, you may have a line that looks like this:

```
# Set the user's default umask
umask 033
```

Be sure to make `root`'s `umask` 077, which will disable read, write, and execute permission for other users, unless explicitly changed using `chmod`. In this case, newly-created directories would have 744 permissions, obtained by subtracting 033 from 777. Newly-created files using the 033 `umask` would have permissions of 644.



In **Mandrake Linux**, it is only necessary to use 002 for a `umask`. This is due to the fact that the default configuration is one user per group.

11.5.2. File Permissions

It's important to ensure that your system files are not open for casual editing by users and groups who shouldn't be doing such system maintenance.

UNIX separates access control on files and directories according to three characteristics: owner, group, and other. There is always exactly one owner, any number of members of the group, and everyone else.

A quick explanation of *UNIX* permissions:

Ownership - Which user(s) and group(s) retain(s) control of the permission settings of the node and parent of the node

Permissions - Bits capable of being set or reset to allow certain types of access to it. Permissions for directories may have a different meaning than the same set of permissions on files.

Read:

- To be able to view contents of a file
- To be able to read a directory

Write:

- To be able to add to or change a file
- To be able to delete or move files in a directory

Execute:

- To be able to run a binary program or *shell* script
- To be able to search in a directory, combined with read permission

Save Text Attribute: (For directories)

The “sticky bit” also has a different meaning when applied to directories than when applied to files. If the sticky bit is set on a directory, then a user may only delete files that he owns or for which he has explicit write permission granted, even when he has write access to the directory. This is designed for directories like /tmp, which are world-writable, but where it may not be desirable to allow any user to delete files at will. The sticky bit is seen as a *t* in a long directory listing.

***suid* Attribute: (For Files)**

This describes set-user-id permissions on the file. When the set user *ID* access mode is set in the owner permissions, and the file is executable, processes which run it are granted access to system resources based on user who owns the file, as opposed to the user who created the process. This is the cause of many “buffer overflow” exploits.

SGID Attribute: (For Files)

If set in the group permissions, this bit controls the “*set group id*” status of a file. This behaves the same way as *suid*, except the group is affected instead. The file must be executable for this to have any effect.

SGID Attribute: (For directories)

If you set the SGID bit on a directory (with `chmod g+s directory`), files created in that directory will have their group set to the directory’s group.

You - The owner of the file

Group - The group you belong to

Everyone - Anyone on the system that is not the owner or a member of the group

File Example:

```
-rw-r--r-- 1 queen users      114 Aug 28 1997 .zlogin
1st bit - directory?          (no)
2nd bit - read by owner?      (yes, by queen)
3rd bit - write by owner?     (yes, by queen)
4th bit - execute by owner?   (no)
5th bit - read by group?      (yes, by users)
6th bit - write by group?     (no)
7th bit - execute by group?   (no)
8th bit - read by everyone?   (yes, by everyone)
9th bit - write by everyone?  (no)
10th bit - execute by everyone? (no)
```

The following lines are examples of the minimum sets of permissions that are required to perform the access described. You may want to give more permission than what’s listed here, but this should describe what these minimum permissions on files do:

```
--r----- Allow read access to the file by owner
--w----- Allows the owner to modify or delete the file
          (Note that anyone with write permission to the directory
          the file is in can overwrite it and thus delete it)
---x----- The owner can execute this program, but not shell scripts,
          which still need read permission
---s----- Will execute with effective User ID = to owner
-----s-   Will execute with effective Group ID = to group
-rw-----T No update of "last modified time". Usually used for swap
          files
---t----- No effect. (formerly sticky bit)
```

Directory Example:

```

drwxr-xr-x  3 queen  users          512 Sep 19 13:47 .public_html/
1st bit - directory?                (yes, it contains many files)
2nd bit - read by owner?            (yes, by queen)
3rd bit - write by owner?           (yes, by queen)
4th bit - execute by owner?         (yes, by queen)
5th bit - read by group?            (yes, by users)
6th bit - write by group?           (no)
7th bit - execute by group?         (yes, by users)
8th bit - read by everyone?         (yes, by everyone)
9th bit - write by everyone?        (no)
10th bit - execute by everyone?     (yes, by everyone)

```

The following lines are examples of the minimum sets of permissions that are required to perform the access described. You may want to give more permission than what's listed, but this should describe what these minimum permissions on directories do:

```

dr----- The contents can be listed, but file attributes can't be read
d--x----- The directory can be entered, and used in full execution
           paths
dr-x----- File attributes can be read by owner
d-wx----- Files can be created/deleted, even if the directory
           isn't the current one
d-----x-t Prevents files from deletion by others with write
           access. Used on /tmp
d---s---s-- No effect

```

System configuration files (usually in /etc) are usually mode 640 (-rw-r-----), and owned by root. Depending on your site's security requirements, you might adjust this. Never leave any system files writable by a group or everyone. Some configuration files, including /etc/shadow, should only be readable by root, and directories in /etc should at least not be accessible by others.

suid shell Scripts

suid shell scripts are a serious security risk, and for this reason the kernel will not honor them. Regardless of how secure you think the *shell* script is, it can be exploited to give the cracker a *root shell*.

11.5.3. Integrity Checking

Another very good way to detect local (and also network) attacks on your system is to run an integrity checker like *Tripwire*, *Aide* or *Osiris*. These integrity checkers run number of checksums on all your important binaries and config files and compares them against a database of former, known-good values as a reference. Thus, any changes in the files will be flagged.

It's a good idea to install these sorts of programs onto a floppy, and then physically set the write protect on the floppy. This way intruders can't tamper with the integrity checker itself or change the database. Once you have something like this setup, it's a good idea to run it as part of your normal security administration duties to see if anything has changed.

You can even add a crontab entry to run the checker from your floppy every night and mail you the results in the morning. Something like:

```

# set mailto
MAILTO=queen
# run Tripwire
15 05 * * * root /usr/local/adm/tcheck/tripwire

```

will mail you a report each morning at 5:15am.

Integrity checkers can be a godsend to detecting intruders before you would otherwise notice them. Since a lot of files change on the average system, you have to be careful what is cracker activity and what is your own doing.

You can find the freely available unsupported version of Tripwire at TripWire (<http://www.tripwire.org>) free of charge. Manuals and support can be purchased.

Aide can be found at <http://www.cs.tut.fi/~rammer/aide.html> (<http://www.cs.tut.fi/~rammer/aide.html>).

Osiris can be found at <http://osiris.shmoo.com/> (<http://www.shmoo.com/osiris/>).

11.5.4. Trojan Horses

“Trojan Horses” are named after the fabled ploy in Homer’s *Iliad*. The idea is that a cracker distributes a program or binary that sounds great, and encourages other people to download it and run it as root. Then the program can compromise their system while they are not paying attention. While they think the binary they just pulled down does one thing (and it might very well), it also compromises their security.

You should take care of what programs you install on your computer. **MandrakeSoft** provides MD5 check-sums and PGP signatures on its RPM files so you can verify you are installing the real thing. You should never run any unfamiliar binary, for which you don’t have the source, as root! Few attackers are willing to release source code to public scrutiny.

Although it can be complex, make sure you are getting the source for a program from its real distribution site. If the program is going to run as root, make sure either you or someone you trust has looked over the source and verified it.

11.6. Password Security and Encryption



Most of encryption programs described in this chapter are available in your **Mandrake Linux** distribution.

One of the most important security features used today are passwords. It is important for both you and all your users to have secure, unguessable passwords. Your **Mandrake Linux** distributions include `passwd` program that do not allow you to set an easy to guess password. Make sure your `passwd` program is up to date.

In-depth discussion of encryption is beyond the scope of this chapter, but an introduction is in order. Encryption is very useful, possibly even necessary in this day and age. There are all sorts of methods of encrypting data, each with its own set of characteristics.

Most *UNIX* systems (and *GNU/Linux* is no exception) primarily use a one-way encryption algorithm, called DES (Data Encryption Standard) to encrypt your passwords. This encrypted password is then stored in `/etc/shadow`. When you attempt to login, the password you type in is encrypted again and compared with the entry in the file that stores your passwords. If they match, it must be the same password, and you are allowed access. Although DES is a two-way encryption algorithm (you can code and then decode a message, given the right keys), the variant that most Unixes use is one-way. This means that it should not be possible to reverse the encryption to get the password from the contents of `/etc/shadow`.

Brute force attacks, such as “Crack” or “John the Ripper” (see Section “Crack” and “John the Ripper”, page 153) can often guess passwords unless your password is sufficiently random. PAM modules (see below) allow you to use a different encryption routine with your passwords (MD5 or the like). You can use Crack to your advantage, as well. Consider periodically running Crack against your own password database, to find insecure passwords. Then contact the offending user, and instruct him to change his password.

You can go to CERN (http://consult.cern.ch/writeup/security/security_3.html) for information on how to choose a good password.

11.6.1. PGP And Public-Key Cryptography

Public-key cryptography, such as that used for PGP, uses one key for encryption, and one key for decryption. Traditional cryptography, however, uses the same key for encryption and decryption; this key must be known to both parties, and thus somehow transferred from one to the other securely.

To alleviate the need to securely transmit the encryption key, public-key encryption uses two separate keys: a public key and a private key. Each person’s public key is available by anyone to do the encryption, while at

the same time each person keeps his or her private key to decrypt messages encrypted with the correct public key.

There are advantages to both public key and private key cryptography, and you can read about those differences in the RSA Cryptography FAQ (<http://www.rsasecurity.com/rsalabs/faq/>), listed at the end of this section.

PGP (Pretty Good Privacy) is well-supported on *GNU/Linux*. Versions 2.6.2 and 5.0 are known to work well. For a good primer on *PGP* and how to use it, take a look at the different *PGP* FAQs at [faqs.org](http://www.faqs.org/faqs/pgp-faq/) (<http://www.faqs.org/faqs/pgp-faq/>)

Be sure to use the version that is applicable to your country. Due to export restrictions by the US Government, strong-encryption is prohibited from being transferred in electronic form outside the country.

US export controls are now managed by EAR (Export Administration Regulations). They are no longer governed by ITAR.

There is also a step-by-step guide for configuring *PGP* on *GNU/Linux* available at LinuxFocus (<http://mercury.chem.pitt.edu/~angel/LinuxFocus/English/November1997/article7.html>). It was written for the international version of *PGP*, but is easily adaptable to the United States version. You may also need a patch for some of the latest versions of *GNU/Linux*; the patch is available at metalab (<ftp://metalab.unc.edu/pub/Linux/apps/crypto>).

There is a project maintaining a free re-implementation of *PGP* with open source. GnuPG is a complete and free replacement for *PGP*. Because it does not use IDEA or RSA it can be used without any restrictions. *GnuPG* is in compliance with OpenPGP (<http://www.faqs.org/rfcs/rfc2440.html>). See the GNU Privacy Guard web page for more information: <http://www.gnupg.org/> (<http://www.gnupg.org/>).

More information on cryptography can be found in the RSA cryptography FAQ, available at rsasecurity (<http://www.rsasecurity.com/rsalabs/faq/>). Here you will find information on such terms as “Diffie-Hellman”, “public-key cryptography”, “digital certificates”, etc.

11.6.2. SSL, S-HTTP and S/MIME

Often users ask about the differences between the various security and encryption protocols, and how to use them. While this isn’t an encryption document, it is a good idea to explain briefly what each protocol is, and where to find more information.

- **SSL:** - SSL, or Secure Sockets Layer, is an encryption method developed by Netscape to provide security over the Internet. It supports several different encryption protocols, and provides client and server authentication. SSL operates at the transport layer, creates a secure encrypted channel of data, and thus can seamlessly encrypt data of many types. This is most commonly seen when going to a secure site to view a secure online document with *Communicator*, and serves as the basis for secure communications with *Communicator*, as well as many other Netscape Communications data encryption. More information can be found at Openssl.org (<http://www.openssl.org>). Information on Netscape’s other security implementations, and a good starting point for these protocols is available at netscape (<http://home.netscape.com/info/security-doc.html>). It’s also worth noting that the SSL protocol can be used to pass many other common protocols, “wrapping” them for security. See quiltaholic (<http://www.quiltaholic.com/rickk/sslwrap/>)
- **S-HTTP:** - S-HTTP is another protocol that provides security services across the Internet. It was designed to provide confidentiality, authentication, integrity, and non-repudiability [cannot be mistaken for someone else] while supporting multiple key-management mechanisms and cryptographic algorithms via option negotiation between the parties involved in each transaction. S-HTTP is limited to the specific software that is implementing it, and encrypts each message individually. [From RSA Cryptography FAQ, page 138]
- **S/MIME:** - S/MIME, or Secure Multipurpose Internet Mail Extension, is an encryption standard used to encrypt electronic mail and other types of messages on the Internet. It is an open standard developed by RSA, so it is likely we will see it on *GNU/Linux* one day soon. More information on S/MIME can be found at RFC2311 (<http://www.ietf.org/rfc/rfc2311.txt>).

11.6.3. IPSEC Implementations

Along with CIPE, and other forms of data encryption, there are also several other implementations of IPSEC for *GNU/Linux*. IPSEC is an effort by the IETF to create cryptographically-secure communications at the IP network level, and to provide authentication, integrity, access control, and confidentiality. Information on IPSEC and Internet draft can be found at <http://www.ietf.org/html.charters/ipsec-charter.html> (<http://www.ietf.org/html.charters/ipsec-charter.html>). You can also find links to other protocols involving key management, and an IPSEC mailing list and archives.

The x-kernel *GNU/Linux* implementation, which was (project is now closed) being developed at the University of Arizona, uses an object-based framework for implementing network protocols called x-kernel, and can be found at <http://www.cs.arizona.edu/xkernel/hpcc-blue/linux.html> (<http://www.cs.arizona.edu/xkernel/hpcc-blue/linux.html>). Most simply, the x-kernel is a method of passing messages at the kernel level, which makes for an easier implementation.

Another freely-available IPSEC implementation is the *GNU/Linux* FreeS/WAN IPSEC. Their web page states, “These services allow you to build secure tunnels through untrusted networks. Everything passing through the untrusted net is encrypted by the IPSEC gateway computer and decrypted by the gateway at the other end. The result is Virtual Private Network or VPN. This is a network which is effectively private even though it includes computers at several different sites connected by the insecure Internet.”

It’s available for download from <http://www.xs4all.nl/~freeswan/> (<http://www.xs4all.nl/~freeswan/>).

As with other forms of cryptography, it is not distributed with the kernel by default due to export restrictions.

11.6.4. ssh (Secure SHell) And stelnet

ssh and stelnet are suites of programs that allow you to login to remote systems and have a encrypted connection.

openssh is a suite of programs used as a secure replacement for rlogin, rsh and rcp. It uses public-key cryptography to encrypt communications between two hosts, as well as to authenticate users. It can be used to securely login to a remote host or copy data between hosts, while preventing man-in-the-middle attacks (session hijacking) and DNS spoofing. It will perform data compression on your connections, and secure X11 communications between hosts.

There are several ssh implementations now. The original commercial implementation by Data Fellows can be found at The ssh home page available at <http://www.datafellows.com> (<http://www.datafellows.com>).

The excellent *Openssh* implementation is based on a early version of the *datafellows* ssh and has been totally reworked to not include any patented or proprietary pieces. It is free and under a BSD license. It can be found at: <http://www.openssh.com> (<http://www.openssh.com>).

There is also a open source project to re-implement ssh from the ground up called “psst...”. For more information see: <http://www.net.lut.ac.uk/psst/> (<http://www.net.lut.ac.uk/psst/>).

You can also use ssh from your *Windows* workstation to your *GNU/Linux* ssh server. There are several freely available *Windows* client implementations, including PuTTY (<http://www.chiark.greenend.org.uk/~sgtatham/putty/>) and the one at therapy ssh (<http://guardian.htu.tuwien.ac.at/therapy/ssh/>) as well as a commercial implementation from DataFellows, at datafellows (<http://www.datafellows.com>).

SSLey (outdated, see OpenSSL below) is a free implementation of Netscape’s Secure Sockets Layer, developed by Eric Young. It includes several applications, such as “Secure telnet”, a module for *Apache*, several databases, as well as several algorithms including DES, IDEA and “Blowfish”.

Using this library, a secure telnet replacement has been created that does encryption over a telnet connection. Unlike SSH, stelnet uses SSL, the Secure Sockets Layer protocol developed by Netscape. You can find Secure telnet and Secure FTP by starting with the SSLey FAQ, available at <http://www.psy.uq.oz.au/~ftp/Crypto/> (<http://www.psy.uq.oz.au/~ftp/Crypto/>).



The *OpenSSL* Project is based on *SSLeasy* and is meant to develop a robust, commercial-grade, full-featured, and Open Source toolkit implementing the Secure Sockets Layer (SSL v2/v3) and Transport Layer Security (TLS v1) protocols as well as a full-strength general purpose cryptography library. For more information about this project, consult the OpenSSL home page (www.openssl.org). There is a large list of applications based on *OpenSSL* at OpenSSL related applications (<http://www.openssl.org/related/apps.html>).

SRP is another secure telnet/ftp implementation. From their web page:

"The SRP project is developing secure Internet software for free worldwide use. Starting with a fully-secure *Telnet* and *FTP* distribution, we hope to supplant weak networked authentication systems with strong replacements that do not sacrifice user-friendliness for security. Security should be the default, not an option!"

For more information, go to stanford.edu (<http://www-cs-students.stanford.edu/~tjw/srp/>).

11.6.5. PAM - Pluggable Authentication Modules

Your version of **Mandrake Linux** distribution ships with a unified authentication scheme called PAM. PAM allows you to change your authentication methods and requirements on the fly, and encapsulate all local authentication methods without recompiling any of your binaries. Configuration of PAM is beyond the scope of this chapter, but be sure to take a look at the PAM web site (<http://www.kernel.org/pub/linux/libs/pam/index.html>) for more information.

Just a few of the things you can do with PAM:

- Use encryption other than DES for your passwords. (Making them harder to brute-force decode)
- Set resource limits on all your users so they can't perform denial-of-service attacks (number of processes, amount of memory, etc.)
- Enable shadow passwords (see below) on the fly
- allow specific users to login only at specific times from specific places

Within a few hours of installing and configuring your system, you can prevent many attacks before they even occur. For example, use PAM to disable the system-wide usage of `.rhosts` files in user's home directories by adding these lines to `/etc/pam.d/rlogin`:

```
#
# Disable rsh/rlogin/rexec for users
#
login auth required pam_rhosts_auth.so no_rhosts
```

11.6.6. Cryptographic IP Encapsulation (CIPE)

The primary goal of this software is to provide a facility for secure (against eavesdropping, including traffic analysis, and faked message injection) subnetwork interconnection across an insecure packet network such as the Internet.

CIPE encrypts the data at the network level. Packets traveling between hosts on the network are encrypted. The encryption engine is placed near the driver which sends and receives packets.

This is unlike SSH, which encrypts the data by connection, at the socket level. A logical connection between programs running on different hosts is encrypted.

CIPE can be used in tunneling, in order to create a Virtual Private Network. Low-level encryption has the advantage that it can be made to work transparently between the two networks connected in the VPN, without any change to application software.

Summarized from the CIPE documentation:

"The IPSEC standards define a set of protocols which can be used (among other things) to build encrypted VPNs. However, IPSEC is a rather heavyweight and complicated protocol set with a lot of options, implementations of the full protocol set are still rarely used and some issues (such as key management) are still not fully resolved. CIPE uses a simpler approach, in which many things which can be parameterized (such as the choice of the actual encryption algorithm used) are an install-time fixed choice. This limits flexibility, but allows for a simple (and therefore efficient, easy to debug...) implementation."

Further information can be found at <http://sites.inka.de/sites/bigred/devel/cipe.html> (<http://sites.inka.de/sites/bigred/devel/cipe.html>)

As with other forms of cryptography, it is not distributed with the kernel by default due to export restrictions.

11.6.7. Kerberos

Kerberos is an authentication system developed by the Athena Project at MIT. When a user logs in, *Kerberos* authenticates that user (using a password), and provides the user with a way to prove her identity to other servers and hosts scattered around the network.

This authentication is then used by programs such as `rlogin` to allow the user to login to other hosts without a password (in place of the `.rhosts` file). This authentication method can also be used by the mail system in order to guarantee that mail is delivered to the correct person, as well as to guarantee that the sender is who he claims to be.

Kerberos and the other programs that come with it, prevent users from "spoofing" the system into believing they are someone else. Unfortunately, installing *Kerberos* is very intrusive, requiring the modification or replacement of numerous standard programs.

You can find more information about Kerberos by looking at the Kerberos FAQ (<http://www.faqs.org/faqs/kerberos-faq/general/>), and the code can be found at <http://web.mit.edu/kerberos/www/> (<http://web.mit.edu/kerberos/www/>).

[From: Stein, Jennifer G., Clifford Neuman, and Jeffrey L. Schiller. "Kerberos: An Authentication Service for Open Network Systems." USENIX Conference Proceedings, Dallas, Texas, Winter 1998.]

Kerberos should not be your first step in improving security of your host. It is quite involved, and not as widely used as, say, SSH.

11.6.8. "Crack" and "John the Ripper"

If for some reason your `passwd` program is not enforcing hard-to-guess passwords, you might want to run a password-cracking program and make sure your users' passwords are secure.

Password cracking programs work on a simple idea: they try every word in the dictionary, and then variations on those words, encrypting each one and checking it against your encrypted password. If they get a match they know what your password is.

There are a number of programs out there...the two most notable of which are *Crack* and *John the Ripper* (See OpenWall (<http://www.openwall.com/john/>)). They will take up a lot of your CPU time, but you should be able to tell if an attacker could get in using them by running them first yourself and notifying users with weak passwords. Note that an attacker would have to use some other hole first in order to read your `/etc/shadow` file, but such holes are more common than you might think.

Because security is only as strong as the most insecure host, it is worth mentioning that if you have any *Windows* computers on your network, you should check out *L0phtCrack*, a *Crack* implementation for *Windows*. It's available from <http://www.atstake.com> (<http://www.atstake.com/research/lc3/>)

11.6.9. CFS – Cryptographic File System And TCFS – Transparent Cryptographic File System

CFS is a way of encrypting an entire directory tree and allowing users to store encrypted files on them. It uses an NFS server running on the local computer. More information and source code are available at [att](ftp://ftp.research.att.com/dist/mab/) (<ftp://ftp.research.att.com/dist/mab/>).

TCFS improves on CFS by adding more integration with the file system, so that it's transparent to users that the file system is encrypted. More information at: [tcfs/](http://www.tcfs.it/) (<http://www.tcfs.it/>).

It also need not be used on entire file systems. It works on directory trees as well.

11.6.10. X11, SVGA And Display Security

11.6.10.1. X11

It's important for you to secure your graphical display to prevent attackers from grabbing your passwords as you type them, reading documents or information you are reading on your screen, or even using a hole to gain root access. Running remote *X* applications over a network also can be fraught with peril, allowing sniffers to see all your interaction with the remote system.

X has a number of access-control mechanisms. The simplest of them is host-based: you use *xhost* to specify the hosts that are allowed access to your display. This is not very secure at all, because if someone has access to your computer, they can *xhost + their computer* and get in easily. Also, if you have to allow access from an untrusted computer, anyone there can compromise your display.

When using *xdm* (*X* Display Manager), or its *KDE* counterpart: *KDM*, to log in, you get a much better access method: MIT-MAGIC-COOKIE-1. A 128-bit "cookie" is generated and stored in your *.Xauthority* file. If you need to allow a remote computer access to your display, you can use the *xauth* command and the information in your *.Xauthority* file to provide access to only that connection. See the Remote-*X*-Apps mini-howto, available at <http://metalab.unc.edu/LDP/HOWTO/mini/Remote-X-Apps.html> (<http://metalab.unc.edu/LDP/HOWTO/mini/Remote-X-Apps.html>).

You can also use *ssh* (see *ssh (Secure SHell) And stelnet*, page 151, above) to allow secure *X* connections. This has the advantage of also being transparent to the end user, and means that no unencrypted data flows across the network.

You can also disable any remote connections to your *X* server by using the *-nolisten tcp* options to your *X* server. This will prevent any network connections to your server over *tcp* sockets.

Take a look at the *Xsecurity* man page for more information on *X* security. The safe bet is to use *xdm* to login to your console and then use *ssh* to go to remote sites on which you wish to run *X* programs.

11.6.10.2. SVGA

SVGAlib programs are typically *suid-root* in order to access all your *GNU/Linux* computer's video hardware. This makes them very dangerous. If they crash, you typically need to reboot your computer to get a usable console back. Make sure any *SVGA* programs you are running are authentic, and can at least be somewhat trusted. Even better, don't run them at all.

11.6.10.3. GGI (Generic Graphics Interface Project)

The *GNU/Linux* GGI project is trying to solve several of the problems with video interfaces on *GNU/Linux*. GGI will move a small piece of the video code into the *GNU/Linux* kernel, and then control access to the video system. This means GGI will be able to restore your console at any time to a known good state. They will also allow a secure attention key, so you can be sure that there is no Trojan horse login program running on your console. <http://www.ggi-project.org/> (<http://www.ggi-project.org/>)

11.7. Kernel Security

This is a description of the kernel configuration options that relate to security, and an explanation of what they do, and how to use them.

As the kernel controls your computer's networking, it is important that it be very secure, and not be compromised. To prevent some of the latest networking attacks, you should try to keep your kernel version current. You can find new kernels at <ftp://ftp.kernel.org> (<ftp://ftp.kernel.org>) or from packages updates available through MandrakeUpdate.

There is also an international group providing a single unified cryptographic patch to the mainstream *GNU/Linux* kernel. This patch provides support for a number of cryptographic subsystems and things that

cannot be included in the mainstream kernel due to export restrictions. For more information, visit their web page at: kerneli (<http://www.kerneli.org>).

11.7.1. Kernel Compile Options

When this document was written, kernel 2.2 was state-of-the-art. Still today, most firewalls still run 2.2. However, with kernel 2.4, a lot of things have changed. Most of the compile options in this chapter are still valid, but the Masquerading and port forwarding have been replaced by iptables. You can have more information on iptables at linuxguruz.org (<http://www.linuxguruz.org/iptables/howto/iptables-HOWTO.html>).

For 2.2.x kernels, the following options apply. You should see these options during the kernel configuration process. Many of the comments here are from `/usr/src/linux/Documentation/Configure.help`, which is the same document that is referenced while using the Help facility during the `make config` stage of compiling the kernel. Please consult the chapter *Compiling And Installing New Kernels* of the *Reference Manual* to a full description of the compilation of a brand new kernel.

- Network Firewalls (CONFIG_FIREWALL)

This option should be on if you intend to run any firewalling or masquerading on your *GNU/Linux* computer. If it's just going to be a regular client computer, it's safe to say no.

- IP: forwarding/gatewaying (CONFIG_IP_FORWARD)

If you enable IP forwarding, your *GNU/Linux* box essentially becomes a router. If your computer is on a network, you could be forwarding data from one network to another, and perhaps subverting a firewall that was put there to prevent this from happening. Normal dial-up users will want to disable this, and other users should concentrate on the security implications of doing this. Firewall computers will want this enabled, and used in conjunction with firewall software.

You can enable IP forwarding dynamically using the following command:

```
root# echo 1 > /proc/sys/net/ipv4/ip_forward
```

and disable it with the command:

```
root# echo 0 > /proc/sys/net/ipv4/ip_forward
```

- IP: syn cookies (CONFIG_SYN_COOKIES)

A "SYN Attack" is a denial of service (DoS) attack that consumes all the resources on your computer, forcing you to reboot. We can't think of a reason you wouldn't normally enable this. In the 2.1 kernel series this config option merely allows syn cookies, but does not enable them. To enable them, you have to do:

```
root# echo 1 > /proc/sys/net/ipv4/tcp_syncookies <P>
```

- IP: Firewalling (CONFIG_IP_FIREWALL)

This option is necessary if you are going to configure your computer as a firewall, do masquerading, or wish to protect your dial-up workstation from someone entering via your PPP dial-up interface.

- IP: firewall packet logging (CONFIG_IP_FIREWALL_VERBOSE)

This option gives you information about packets your firewall received, like sender, recipient, port, etc.

- IP: Drop source routed frames (CONFIG_IP_NOSR)

This option should be enabled. Source routed frames contain the entire path to their destination inside of the packet. This means that routers through which the packet goes do not need to inspect it, and just forward it on. This could lead to data entering your system that may be a potential exploit.

- IP: masquerading (CONFIG_IP_MASQUERADE)

If one of the computers on your local network for which your *GNU/Linux* box acts as a firewall wants to send something to the outside, your box can “masquerade” as that host, i.e., it forwards the traffic to the intended destination, but makes it look like it came from the firewall box itself. See <http://www.indyramp.com/masq> (<http://www.indyramp.com/masq>) and the chapter “Configuring Masqueraded Clients”, page 101 for more information.

- IP: ICMP masquerading (CONFIG_IP_MASQUERADE_ICMP)

This option adds ICMP masquerading to the previous option of only masquerading TCP or UDP traffic.

- IP: transparent proxy support (CONFIG_IP_TRANSPARENT_PROXY)

This enables your *GNU/Linux* firewall to transparently redirect any network traffic originating from the local network and destined for a remote host to a local server, called a “transparent proxy server”. This makes the local computers think they are talking to the remote end, while in fact they are connected to the local proxy. See the IP-Masquerading *HOWTO* and <http://www.indyramp.com/masq> (<http://www.indyramp.com/masq>) for more information.

- IP: always defragment (CONFIG_IP_ALWAYS_DEFRAG)

Generally this option is disabled, but if you are building a firewall or a masquerading host, you will want to enable it. When data is sent from one host to another, it does not always get sent as a single packet of data, but rather it is fragmented into several pieces. The problem with this is that the port numbers are only stored in the first fragment. This means that someone can insert information into the remaining packets that isn’t supposed to be there. It could also prevent a teardrop attack against an internal host that is not yet itself patched against it.

- Packet Signatures (CONFIG_NCPFS_PACKET_SIGNING)

This is an option that will sign NCP packets for stronger security. Normally you can leave it off, but it is there if you do need it.

- IP: Firewall packet netlink device (CONFIG_IP_FIREWALL_NETLINK)

This is a really neat option that allows you to analyze the first 128 bytes of the packets in a user-space program, to determine if you would like to accept or deny the packet, based on its validity.

- Socket Filtering (CONFIG_FILTER)

For most people, it’s safe to say no to this option. This option allows you to connect a user-space filter to any socket and determine if packets should be allowed or denied. Unless you have a very specific need and are capable of programming such a filter, you should say no. Also note that as of this writing, all protocols were supported except TCP.

- Port Forwarding

Port Forwarding is an addition to IP Masquerading which allows some forwarding of packets from outside to inside a firewall on given ports. This could be useful if, for example, you want to run a web server behind the firewall or masquerading host and that web server should be accessible from the outside world. An external client sends a request to port 80 of the firewall, the firewall forwards this request to the web server, the web server handles the request and the results are sent through the firewall to the original client. The client thinks that the firewall computer itself is running the web server. This can also be used for load balancing if you have a farm of identical web servers behind the firewall. Information about this feature is available from monmouth (<http://www.monmouth.demon.co.uk/ipsubs/portforwarding.html>). For general info, please see compsoc (<ftp://ftp.compsoc.net/users/steve/ipportfw/linux21/>).

- Socket Filtering (CONFIG_FILTER)

Using this option, user-space programs can attach a filter to any socket and thereby tell the kernel that it should allow or disallow certain types of data to get through the socket. *GNU/Linux* socket filtering works on all socket types except TCP for now. See the text file `/usr/src/linux/Documentation/networking/filter.txt` for more information.

- IP: Masquerading

The 2.2 kernel masquerading has been improved. It provides additional support for masquerading special protocols, etc. Be sure to read the IP Chains *HOWTO* for more information.

11.7.2. Kernel Devices

There are a few block and character devices available on *GNU/Linux* that will also help you with security.

The two devices `/dev/random` and `/dev/urandom` are provided by the kernel to provide random data at any time.

Both `/dev/random` and `/dev/urandom` should be secure enough to use in generating *PGP* keys, *ssh* challenges, and other applications where secure random numbers are required. Attackers should be unable to predict the next number given any initial sequence of numbers from these sources. There has been a lot of effort put in to ensuring that the numbers you get from these sources are random in every sense of the word.

The only difference between the two devices, is that `/dev/random` runs out of random bytes and it makes you wait for more to be accumulated. Note that on some systems, it can block for a long time waiting for new user-generated entropy to be entered into the system. So you have to use care before using `/dev/random`. (Perhaps the best thing to do is to use it when you're generating sensitive keying information, and you tell the user to pound on the keyboard repeatedly until you print out "OK, enough".)

`/dev/random` is high quality entropy, generated from measuring the inter-interrupt times etc. It blocks until enough bits of random data are available.

`/dev/urandom` is similar, but when the store of entropy is running low, it'll return a cryptographically strong hash of what there is. This isn't as secure, but it's enough for most applications.

You might read from the devices using something like:

```
root# head -c 6 /dev/urandom | mimencode
```

This will print six random characters on the console, suitable for password generation. You can find *mimencode* in the *metamail* package.

See `/usr/src/linux/drivers/char/random.c` for a description of the algorithm.

11.8. Network Security

Network security is becoming more and more important as people spend more and more time connected. Compromising network security is often much easier than compromising physical or local security, and is much more common.

There are a number of good tools to assist with network security, and more and more of them are shipped with your **Mandrake Linux** distribution, either in the main CD-ROM, contribs, or through the FTP crypto server (see above).

11.8.1. Packet Sniffers

One of the most common ways intruders gain access to more systems on your network is by employing a packet sniffer on a already compromised host. This "sniffer" just listens on the *Ethernet* port for things like `passwd` and `login` and `su` in the packet stream and then logs the traffic after that. This way, attackers gain passwords for systems they are not even attempting to break into. Clear-text passwords are very vulnerable to this attack.

Example: Host A has been compromised. Attacker installs a sniffer. Sniffer picks up admin logging into Host B from Host C. It gets the admin's personal password as they login to B. Then, the admin does a `su` to fix a problem. They now have the root password for Host B. Later the admin lets someone `telnet` from his account to Host Z on another site. Now the attacker has a password/*login* on Host Z.

In this day and age, the attacker doesn't even need to compromise a system to do this: they could also bring a laptop or PC into a building and tap into your net.

Using `ssh` or other encrypted password methods thwarts this attack. Things like APOP for POP accounts also prevents this attack. (Normal POP logins are very vulnerable to this, as is anything that sends clear-text passwords over the network.)

11.8.2. System Services and `tcp_wrappers`

Before you put your *GNU/Linux* system on **ANY** network the first thing to look at is what services you need to offer. Services that you do not need to offer should be disabled so that you have one less thing to worry about and attackers have one less place to look for a hole.

There are a number of ways to disable services under *GNU/Linux*. You can look at your `/etc/inetd.conf` file and see what services are being offered by your `inetd`. Disable any that you do not need by commenting them out (# at the beginning of the line), and then restart your `inetd` service.

You can also remove (or comment out) services in your `/etc/services` file. This will mean that local clients will also be unable to find the service (i.e., if you remove `ftp`, and try and `ftp` to a remote site from that computer it will fail with an `unknown service` message). It's usually not worth the trouble to remove services from `/etc/services`, since it provides no additional security. If a local person wanted to use `ftp` even though you had commented it out, they would make their own client that use the common FTP port and would still work fine.

Some of the services you might want to leave enabled are:

- `ftp`
- `telnet` (or `ssh`)
- mail, such as `pop-3` or `imap`
- `identd`

If you know you are not going to use some particular package, you can also delete it entirely. `rpm -e packagename` will erase an entire package.

Additionally, you really want to disable the `rsh`/`rlogin`/`rcp` utilities, including `login` (used by `rlogin`), `shell` (used by `rcp`), and `exec` (used by `rsh`) from being started in `/etc/inetd.conf`. These protocols are extremely insecure and have been the cause of exploits in the past.

You should check your `/etc/rc.d/rc[0-9].d`, and see if any of the servers started in that directory are not needed. The files in that directory are actually symbolic links to files in the directory `/etc/rc.d/init.d`. Renaming the files in the `init.d` directory disables all the symbolic links that point to that file. If you only wish to disable a service for a particular run level, rename the appropriate symbolic link by replacing the `S` with a `K`, like this:

```
root# cd /etc/rc6.d
root# mv S45dhcpd K45dhcpd
```



You may also use a command line utility to do that: `chkconfig` or the graphical interface under *KDE*: `kysysv`.

Your **Mandrake Linux** distributions ships with a `tcp_wrapper` “wrapping” all your TCP services. The `tcp_wrapper` (`tcpd`) is invoked from `inetd` instead of the real server. `tcpd` then checks the host that is requesting the service, and either executes the real server, or denies access from that host. `tcpd` allows you to restrict access to your TCP services. You should edit `/etc/hosts.allow` and add in only those hosts that need to have access to your computer’s services.

If you are a home dial up user, we suggest you deny ALL. `tcpd` also logs failed attempts to access services, so this can alert you if you are under attack. If you add new services, you should be sure to configure them to use `tcp_wrappers` if they are TCP-based. For example, a normal dial-up user can prevent outsiders from connecting to his computer, yet still have the ability to retrieve mail, and make network connections to the Internet. To do this, you might add the following to your `/etc/hosts.allow`:

ALL: 127.

And of course `/etc/hosts.deny` would contain:

ALL: ALL

which will prevent external connections to your computer, yet still allow you from the inside to connect to servers on the Internet.

Keep in mind that `tcp_wrappers` only protects services executed from `inetd`, and a select few others. There very well may be other services running on your computer. You can use `netstat -ta` to find a list of all the services your computer is offering.

11.8.3. Verify Your DNS Information

Keeping up-to-date DNS information about all hosts on your network can help to increase security. If an unauthorized host becomes connected to your network, you can recognize it by its lack of a DNS entry. Many services can be configured to not accept connections from hosts that do not have valid DNS entries.

11.8.4. `identd`

`identd` is a small program that typically runs out of your `inetd` server. It keeps track of what user is running what TCP service, and then reports this to whoever requests it.

Many people misunderstand the usefulness of `identd`, and so disable it or block all off site requests for it. `identd` is not there to help out remote sites. There is no way of knowing if the data you get from the remote `identd` is correct or not. There is no authentication in `identd` requests.

Why would you want to run it then? Because it helps **you** out, and is another data-point in tracking. If your `identd` is not compromised, then you know it’s telling remote sites the user-name or UID of people using TCP services. If the admin at a remote site comes back to you and tells you user so-and-so was trying to hack into their site, you can easily take action against that user. If you are not running `identd`, you will have to look at lots and lots of logs, figure out who was on at the time, and in general take a lot more time to track down the user.

The `identd` that ships with most distributions is more configurable than many people think. You can disable it for specific users (they can make a `.noident` file), you can log all `identd` requests (We recommend it), you can even have `identd` return a UID instead of a user name or even `NO-USER`.

11.8.5. Configuring And Securing The Postfix MTA

The Postfix mail server was written by Wietse Venema, author of Postfix and several other staple Internet security products, as an “attempt to provide an alternative to the widely-used Sendmail program. Postfix attempts to be fast, easy to administer, and hopefully secure, while at the same time being Sendmail-compatible enough to not upset your users.”

Further information on postfix can be found at the Postfix home (<http://www.postfix.org>) and in the Configuring and Securing Postfix (http://www.linuxsecurity.com/feature_stories/feature_story-91.html).

11.8.6. SATAN, ISS, And Other Network Scanners

There are a number of different software packages out there that do port and service-based scanning of computers or networks. *SATAN*, *ISS*, *SAINT*, and *Nessus* are some of the more well-known ones. This software connects to the target computer (or all the target computers on a network) on all the ports they can, and try to determine what service is running there. Based on this information, you can tell if the computer is vulnerable to a specific exploit on that server.

SATAN (Security Administrator's Tool for Analyzing Networks) is a port scanner with a web interface. It can be configured to do light, medium, or strong checks on a computer or a network of computers. It's a good idea to get *SATAN* and scan your computer or network, and fix the problems it finds. Make sure you get the copy of *SATAN* from metalab (<http://metalab.unc.edu/pub/packages/security/Satan-for-Linux/>) or a reputable FTP or web site. There was a Trojan copy of *SATAN* that was distributed out on the net. trouble.org (<http://www.trouble.org/~zen/satan/satan.html>). Note that *SATAN* has not been updated in quite a while, and some of the other tools below might do a better job.

ISS (Internet Security Scanner) is another port-based scanner. It is faster than *Satan*, and thus might be better for large networks. However, *SATAN* tends to provide more information.

TriSentry (formerly Abacus) is a suite of tools to provide host-based security and intrusion detection. look at its home page on the web for more information. <http://www.psionic.com/products/> (<http://www.psionic.com/products/>)

SAINT is an updated version of *SATAN*. It is web based and has many more up to date tests than *SATAN*. You can find out more about it at: <http://www.wwdsi.com/~saint> (<http://www.wwdsi.com/saint>)

Nessus is a free security scanner. It has a GTK graphical interface for ease of use. It is also designed with a very nice plug in setup for new port-scanning tests. For more information, take a look at: <http://www.nessus.org> (<http://www.nessus.org/>)

11.8.6.1. Detecting Port Scans

There are some tools designed to alert you to probes by *SATAN* and *ISS* and other scanning software. However, if you liberally use *tcp_wrappers*, look over your log files regularly, you should be able to notice such probes. Even on the lowest setting, *SATAN* still leaves traces in the logs.

There are also "stealth" port scanners. A packet with the TCP ACK bit set (as is done with established connections) will likely get through a packet-filtering firewall. The returned RST packet from a port that **had no established session** can be taken as proof of life on that port. I don't think TCP wrappers will detect this.

You might also look at SNORT (<http://www.snort.org>), which is a free IDS (Intrusion Detection System), which can detect other network intrusions.

11.8.7. Sendmail, qmail and MTA's¹

One of the most important services you can provide is a mail server. Unfortunately, it is also one of the most vulnerable to attack, simply due to the number of tasks it must perform and the privileges it typically needs.

If you are using *sendmail* it is very important to keep up on current versions. *sendmail* has a long long history of security exploits. Always make sure you are running the most recent version from *sendmail* (<http://www.sendmail.org/>).

Keep in mind that *sendmail* does not have to be running in order for you to send mail. If you are a home user, you can disable *sendmail* entirely, and simply use your mail client to send mail. You might also choose to remove the *-bd* flag from the *sendmail* startup file, thereby disabling incoming requests for mail. In other words, you can execute *sendmail* from your startup script using the following instead:

```
# /usr/lib/sendmail -q15m
```

This will cause *sendmail* to flush the mail queue every fifteen minutes for any messages that could not be successfully delivered on the first attempt.

1. Mail Transport Agents

Many administrators choose not to use *sendmail*, and instead choose one of the other mail transport agents. You might consider switching over to *qmail*. *qmail* was designed with security in mind from the ground up. It's fast, stable, and secure. *Qmail* can be found at *qmail* (<http://www.qmail.org>)

In direct competition to *qmail* is *postfix*, written by Wietse Venema, the author of *tcp_wrappers* and other security tools. Formerly called *vmailer*, and sponsored by **IBM**, this is also a mail transport agent written from the ground up with security in mind. You can find more information about *postfix* at *postfix* (<http://www.postfix.org>)



postfix is the default MTA shipped with **Mandrake Linux**.

11.8.8. Denial of Service (DoS) Attacks

A “Denial of Service” (DoS) attack is one where the attacker tries to make some resource too busy to answer legitimate requests, or to deny legitimate users access to your computer.

Denial of service attacks have increased greatly in recent years. Some of the more popular and recent ones are listed below. Note that new ones show up all the time, so this is just a few examples. Read the *GNU/Linux* security lists and the bugtraq list and archives for more current information.

- **SYN Flooding** - SYN flooding is a network denial of service attack. It takes advantage of a “loophole” in the way TCP connections are created. The newer *GNU/Linux* kernels (2.0.30 and up) have several configurable options to prevent SYN flood attacks from denying people access to your computer or services. See *Kernel Security*, page 154 for proper kernel protection options.
- **Ping Flooding** - Ping flooding is a simple brute-force denial of service attack. The attacker sends a “flood” of ICMP packets to your computer. If they are doing this from a host with better bandwidth than yours, your computer will be unable to send anything on the network. A variation on this attack, called “smurfing”, sends ICMP packets to a host with **your** computer’s return IP, allowing them to flood you less detectably. You can find more information about the “smurf” attack at *linuxsecurity.com* (http://www.linuxsecurity.com/articles/network_security_article-4258.html)

If you are ever under a ping flood attack, use a tool like *tcpdump* to determine where the packets are coming from (or appear to be coming from), then contact your provider with this information. Ping floods can most easily be stopped at the router level or by using a firewall.

- **Ping o’ Death** - The Ping o’ Death attack sends ICMP ECHO REQUEST packets that are too large to fit in the kernel data structures intended to store them. Because sending a single, large (65,510 bytes) “ping” packet to many systems will cause them to hang or even crash, this problem was quickly dubbed the “Ping o’ Death”. This one has long been fixed, and is no longer anything to worry about.

You can find code for most exploits, and a more in-depth description of how they work, at *insecure.org* (<http://www.insecure.org/sploits.html>) using their search engine.

11.8.9. NFS (Network File System) Security

NFS is a very widely-used file sharing protocol. It allows servers running *nfsd* and *mountd* to “export” entire file systems to other computers using NFS file system support built in to their kernels (or some other client support if they are not *GNU/Linux* computers). *mountd* keeps track of mounted file systems in */etc/mtab*, and can display them with *showmount*.

Many sites use NFS to serve home directories to users, so that no matter what computer in the cluster they login to, they will have all their home files.

There is some small amount of security allowed in exporting file systems. You can make your *nfsd* map the remote root user (UID=0) to the *nobody* user, denying them total access to the files exported. However, since individual users have access to their own (or at least the same UID) files, the remote root user can login or *su* to their account and have total access to their files. This is only a small hindrance to an attacker that has access to mount your remote file systems.

If you must use NFS, make sure you export to only those computers that you really need to. Never export your entire root directory; export only directories you need to export.

See the NFS *HOWTO* for more information on NFS, available at LDP (<http://www.ibiblio.org/mdw/HOWTO/NFS-HOWTO/>)

11.8.10. NIS (Network Information Service)

Network Information Service (formerly YP) is a means of distributing information to a group of computers. The NIS master holds the information tables and converts them into NIS map files. These maps are then served over the network, allowing NIS client computers to get login, password, home directory and *shell* information (all the information in a standard `/etc/passwd` file). This allows users to change their password once and have it take effect on all the computers in the NIS domain.

NIS is not at all secure. It was never meant to be. It was meant to be handy and useful. Anyone that can guess the name of your NIS domain (anywhere on the net) can get a copy of your `passwd` file, and use *crack* and *John the Ripper* against your users' passwords. Also, it is possible to spoof NIS and do all sorts of nasty tricks. If you must use NIS, make sure you are aware of the dangers.

There is a much more secure replacement for NIS, called *NIS+*. Check out the NIS *HOWTO* for more information: NIS-HOWTO (<http://www.ibiblio.org/mdw/HOWTO/NIS-HOWTO/>)

11.8.11. Firewalls

Firewalls are a means of controlling what information is allowed into and out of your local network. Typically the firewall host is connected to the Internet and your local LAN, and the only access from your LAN to the Internet is through the firewall. This way the firewall can control what passes back and forth from the Internet and your LAN.

There are a number of types of firewalls and methods of setting them up. *GNU/Linux* computers make pretty good firewalls. Firewall code can be built right into 2.0 and higher kernels. The user-space tools *ipchains* for 2.2 kernels, and *iptables* for 2.4 kernels, allows you to change, on the fly, the types of network traffic you allow. You can also log particular types of network traffic.

Firewalls are a very useful and important technique in securing your network. However, never think that because you have a firewall, you don't need to secure the computers behind it. This is a fatal mistake. Check out the very good Firewall-HOWTO at your latest metalab archive for more information on firewalls and *GNU/Linux*. Firewall-HOWTO (<http://www.ibiblio.org/mdw/HOWTO/Firewall-HOWTO.html>)

If you have no experience with firewalls, and plan to set up one for more than just a simple security policy, the *Firewalls* book by O'Reilly and Associates or other online firewall document is mandatory reading. Check out O'Reilly site (<http://www.ora.com>) for more information. The National Institute of Standards and Technology have put together an excellent document on firewalls. Although dated 1995, it is still quite good. You can find it at nist.gov (<http://cs-www.ncsl.nist.gov/publications/nistpubs/800-10/main.html>). Also of interest includes:

- The Freefire Project -- a list of freely-available firewall tools, available at freefire (http://sites.inka.de/sites/lina/freefire-1/index_en.html)
- *Mason* -- the automated firewall builder for *GNU/Linux*. This is a firewall script that learns as you do the things you need to do on your network! More info at: mason (<http://www.pobox.com/~wstearns/mason/>)

11.8.12. IP Chains – GNU/Linux Kernel 2.2.x Firewalling

GNU/Linux IP Firewalling Chains is an update to the 2.0 *GNU/Linux* firewalling code for the 2.2 kernel. It has many more features than previous implementations, including:

- More flexible packet manipulations
- More complex accounting

- Simple policy changes possible automatically
- Fragments can be explicitly blocked, denied, etc.
- Logs suspicious packets.
- Can handle protocols other than ICMP/TCP/UDP.

Be sure to read the IP Chains *HOWTO* for further information. It is available at TLDP (<http://www.tldp.org/HOWTO/IPCHAINS-HOWTO.html>)

11.8.13. Netfilter – Linux Kernel 2.4.x Firewalling

In yet another set of advancements to the kernel IP packet filtering code, netfilter allows users to set up, maintain, and inspect the packet filtering rules in the new 2.4 kernel.

The netfilter subsystem is a complete rewrite of previous packet filtering implementations including ipchains and ipfwadm. Netfilter provides a large number of improvements, and it has now become an even more mature and robust solution for protecting corporate networks.

iptables is the command-line interface used to manipulate the firewall tables within the kernel.

Netfilter provides a raw framework for manipulating packets as they traverse through various parts of the kernel. Part of this framework includes support for masquerading, standard packet filtering, and now more complete network address translation. It even includes improved support for load balancing requests for a particular service among a group of servers behind the firewall.

The stateful inspection features are especially powerful. Stateful inspection provides the ability to track and control the flow of communication passing through the filter. The ability to keep track of state and context information about a session not only makes rules simpler but also helps to better interpret higher-level protocols.

Additionally, small modules can be developed to perform additional specific functions, such as passing packets to programs in userspace for processing then reinjecting back into the normal packet flow. The ability to develop these programs in userspace reduces the level of complexity that was previously associated with having to make changes directly at the kernel level.

Other IP Tables references include:

- Oskar Andreasson IP Tables Tutorial (http://www.linuxsecurity.com/feature_stories/feature_story-94.html) — Oskar Andreasson speaks with LinuxSecurity.com about his comprehensive IP Tables tutorial and how this document can be used to build a robust firewall for your organization.
- Hal Burgiss Introduces Linux Security Quick-Start Guides (http://www.linuxsecurity.com/feature_stories/feature_story-93.html) — Hal Burgiss has written two authoritative guides on securing Linux, including managing firewalling.
- Netfilter Home page (<http://netfilter.samba.org>) — The netfilter/iptables home page.
- Linux Kernel 2.4 Firewalling Matures: netfilter (http://www.linuxsecurity.com/feature_stories/kernel-netfilter.html) — This LinuxSecurity.com article describes the basics of packet filtering, how to get started using iptables, and a list of the new features available in the latest generation of firewalling for Linux.

11.8.14. VPNs – Virtual Private Networks

VPN's are a way to establish a “virtual” network on top of some already existing network. This virtual network often is encrypted and passes traffic only to and from some known entities that have joined the network. VPN's are often used to connect someone working at home over the public Internet to a internal company network.

If you are running a *GNU/Linux* masquerading firewall and need to pass MS PPTP (Microsoft's VPN point-to-point product) packets, there is a linux kernel patch out to do just that. See: ip-masq-vpn (ftp://ftp.rubyriver.com/pub/jhardin/masquerade/ip_masq_vpn.html).

There are several *GNU/Linux* VPN solutions available:

- `vpnd`. See the <http://sunsite.auc.dk/vpnd/> (<http://sunsite.auc.dk/vpnd/>).
- Free S/Wan, available at <http://www.xs4all.nl/~freeswan/> (<http://www.xs4all.nl/~freeswan/>)
- `ssh` can be used to construct a VPN. See the VPN mini-howto for more information.
- `ups` (virtual private server) at strongcrypto (<http://www.strongcrypto.com>).
- `vtun` (virtual tunnel) at sourceforge (<http://vtun.sourceforge.net/>).
- `yavipin` (<http://yavipin.sourceforge.net>).

See also the section on IPSEC for pointers and more information.

11.9. Security Preparation (Before You Go On-Line)

Ok, so you have checked over your system, and determined it's as secure as feasible, and you're ready to put it online. There are a few things you should now do in order to prepare for an intrusion, so you can quickly disable the intruder, and get back up and running.

11.9.1. Make a Full Backup of Your Computer

Discussion of backup methods and storage is beyond the scope of this chapter, but here are a few words relating to backups and security:

If you have less than 650MB of data to store on a partition, a CD-R copy of your data is a good way to go (as it's hard to tamper with later, and if stored properly can last a long time), you will of course need at least 650MB of space to make the image. Tapes and other re-writable media should be write-protected as soon as your backup is complete, and then verified to prevent tampering. Make sure you store your backups in a secure off-line area. A good backup will ensure that you have a known good point to restore your system from.

11.9.2. Choosing a Good Backup Schedule

A six-tape cycle is easy to maintain. This includes four tapes for during the week, one tape for even Fridays, and one tape for odd Fridays. Perform an incremental backup every day, and a full backup on the appropriate Friday tape. If you make some particularly important changes or add some important data to your system, a full backup might well be in order.

11.9.3. Testing Your Backups

You should do periodic tests of your backups to make sure they are working as you might expect them to. Restores of files and checking against the real data, sizes and listings of backups, and reading old backups should be done on a regular basis.

11.9.4. Backup Your RPM File Database

In the event of an intrusion, you can use your RPM database like you would use `tripwire`, but only if you can be sure it too hasn't been modified. You should copy the RPM database to a floppy, and keep this copy off-line at all times.

The files `/var/lib/rpm/fileindex.rpm` and `/var/lib/rpm/packages.rpm` most likely won't fit on a single floppy. But if compressed, each should fit on a separate floppy.

Now, when your system is compromised, you can use the command:

```
root# rpm -Va
```

to verify each file on the system. See the `rpm` man page, as there are a few other options that can be included to make it less verbose. Keep in mind you must also be sure your RPM binary has not been compromised.

This means that every time a new RPM is added to the system, the RPM database will need to be re-archived. You will have to decide the advantages versus drawbacks.

11.9.5. Keep Track of your System Accounting Data

It is very important that the information that comes from `syslog` has not been compromised. Making the files in `/var/log` readable and writable by only a limited number of users is a good start.

Be sure to keep an eye on what gets written there, especially under the `auth` facility. Multiple login failures, for example, can indicate an attempted break-in.

You will want to look in `/var/log` and check `messages`, `mail.log`, and others.

You might also want to configure your log-rotating script to keep logs around longer so you have time to examine them. Take a look at the `logrotate` man page.

If your log files have been tampered with, see if you can determine when the tampering started, and what sort of things appeared to be tampered with. Are there large periods of time that cannot be accounted for? Checking backup tapes (if you have any) for untampered log files is a good idea.

Intruders typically modify log files in order to cover their tracks, but they should still be checked for strange happenings. You may notice the intruder attempting to gain entrance, or exploit a program in order to obtain the root account. You might see log entries before the intruder has time to modify them.

You should also be sure to separate the `auth` facility from other log data, including attempts to switch users using `su`, login attempts, and other user accounting information.

If possible, configure `syslog` to send a copy of the most important data to a secure system. This will prevent an intruder from covering his tracks by deleting his `login/su/ftp` etc attempts. See the `syslog.conf` man page, and refer to the `@` option.

There are several more advanced `syslogd` programs out there. Take a look at <http://www.core-sdi.com/ssyslog/> (<http://www.core-sdi.com/ssyslog/>) for Secure Syslog. Secure Syslog allows you to encrypt your `syslog` entries and make sure no one has tampered with them.

Another `syslogd` with more features is `syslog-ng` (<http://www.balabit.hu/en/downloads/syslog-ng/>). It allows you a lot more flexibility in your logging and also can crypt your remote `syslog` streams to prevent tampering.

Finally, log files are much less useful when no one is reading them. Take some time out every once in a while to look over your log files, and get a feeling for what they look like on a normal day. Knowing this can help make unusual things stand out.

11.9.6. Apply All New System Updates

Due to the fast-paced nature of security fixes, new (fixed) programs are always being released. Before you connect your computer to the network, it's a good idea to run `MandrakeUpdate` (on another computer connected to the Internet and get all the updated packages since you received your distribution CD-ROM. Many times these packages contain important security fixes, so it's a good idea to get them installed.

11.10. What to Do During and After a Breaking

So you have followed some of the advice here (or elsewhere) and have detected a break-in? The first thing to do is to remain calm. Hasty actions can cause more harm than the attacker would have.

11.10.1. Security Compromise Underway

Spotting a security compromise under way can be a tense undertaking. How you react can have large consequences.

If the compromise you are seeing is a physical one, odds are you have spotted someone who has broken into your home, office or lab. You should notify your local authorities. In a lab, you might have spotted someone trying to open a case or reboot a computer. Depending on your authority and procedures, you might ask them to stop, or contact your local security people.

If you have detected a local user trying to compromise your security, the first thing to do is confirm they are in fact who you think they are. Check the site they are logging in from. Is it the site they normally log in from? No? Then use a non-electronic means of getting in touch. For instance, call them on the phone or walk over to their office/house and talk to them. If they agree that they are on, you can ask them to explain what they

were doing or tell them to cease doing it. If they are not on, and have no idea what you are talking about, odds are this incident requires further investigation. Look into such incidents, and have lots of information before making any accusations.

If you have detected a network compromise, the first thing to do (if you are able) is to disconnect your network. If they are connected via modem, unplug the modem cable; if they are connected via *Ethernet*, unplug the *Ethernet* cable. This will prevent them from doing any further damage, and they will probably see it as a network problem rather than detection.

If you are unable to disconnect the network (if you have a busy site, or you do not have physical control of your computers), the next best step is to use something like `tcp_wrappers` or `ipfwadm` to deny access from the intruder's site.

If you can't deny all people from the same site as the intruder, locking the user's account will have to do. Note that locking an account is not an easy thing. You have to keep in mind `.rhosts` files, FTP access, and a host of possible backdoors.

After you have done one of the above (disconnected the network, denied access from their site, and/or disabled their account), you need to kill all their user processes and log them off.

You should monitor your site well for the next few minutes, as the attacker will try to get back in. Perhaps using a different account, and/or from a different network address.

11.10.2. Security Compromise Has Already Happened

So you have either detected a compromise that has already happened or you have detected it and locked (hopefully) the offending attacker out of your system. Now what?

11.10.2.1. Closing the Hole

If you are able to determine what means the attacker used to get into your system, you should try to close that hole. For instance, perhaps you see several FTP entries just before the user logged in. Disable the FTP service and check and see if there is an updated version, or if any of the lists know of a fix.

Check all your log files, and make a visit to your security lists and pages and see if there are any new common exploits you can fix. You can find your **Mandrake Linux** security fixes by running the `MandrakeUpdate` regularly.

There is now a *GNU/Linux* security auditing project. They are methodically going through all the user-space utilities and looking for possible security exploits and overflows. From their announcement:

"We are attempting a systematic audit of *GNU/Linux* sources with a view to being as secure as *OpenBSD*. We have already uncovered (and fixed) some problems, but more help is welcome. The list is unmoderated and also a useful resource for general security discussions. The list address is: `security-audit@ferret.lmh.ox.ac.uk`. To subscribe, send a mail to: `security-audit-subscribe@ferret.lmh.ox.ac.uk`"

If you don't lock the attacker out, they will likely be back. Not just back on your computer, but back somewhere on your network. If they were running a packet sniffer, odds are good they have access to other local computers.

11.10.2.2. Assessing the Damage

The first thing is to assess the damage. What has been compromised? If you are running an Integrity Checker like *Tripwire*, you can use it to perform an integrity check; and it should help to tell you what has been compromised. If not, you will have to look around at all your important data.

Since *GNU/Linux* systems are getting easier and easier to install, you might consider saving your config files, wiping your disk(s), reinstalling, then restoring your user files and your config files from backups. This will ensure that you have a new, clean system. If you have to backup files from the compromised system, be especially cautious of any binaries that you restore, as they may be Trojan horses placed there by the intruder.

Re-installation should be considered mandatory upon an intruder obtaining root access. Additionally, you'd like to keep any evidence there is, so having a spare disk in the safe may make sense.

Then you have to worry about how long ago the compromise happened, and whether the backups hold any damaged work. More on backups later.

11.10.2.3. Backups, Backups, Backups!

Having regular backups is a godsend for security matters. If your system is compromised, you can restore the data you need from backups. Of course, some data is valuable to the attacker too, and they will not only destroy it, they will steal it and have their own copies; but at least you will still have the data.

You should check several backups back into the past before restoring a file that has been tampered with. The intruder could have compromised your files long ago, and you could have made many successful backups of the compromised file!

Of course, there are also a raft of security concerns with backups. Make sure you are storing them in a secure place. Know who has access to them. (If an attacker can get your backups, they can have access to all your data without you ever knowing it.)

11.10.2.4. Tracking Down the Intruder

Ok, you have locked the intruder out, and recovered your system, but you're not quite done yet. While it is unlikely that most intruders will ever be caught, you should report the attack.

You should report the attack to the admin contact at the site from which the attacker attacked your system. You can look up this contact with `whois` or the Internic database. You might send them an email with all applicable log entries and dates and times. If you spotted anything else distinctive about your intruder, you might mention that too. After sending the email, you should (if you are so inclined) follow up with a phone call. If that admin in turn spots your attacker, they might be able to talk to the admin of the site where they are coming from and so on.

Good crackers often use many intermediate systems, some (or many) of which may not even know they have been compromised. Trying to track a cracker back to their home system can be difficult. Being polite to the admins you talk to can go a long way to getting help from them.

You should also notify any security organizations you are a part of (CERT (<http://www.cert.org/>) or similar), as well as **MandrakeSoft** (<http://www.mandrakesecure.net/en/>).

11.11. Security Sources

There are a **lot** of good sites out there for *UNIX* security in general and *GNU/Linux* security specifically. It's very important to subscribe to one (or more) of the security mailing lists and keep current on security fixes. Most of these lists are very low volume, and very informative.

11.11.1. LinuxSecurity.com References

The LinuxSecurity.com web site has numerous Linux and open source security references written by the LinuxSecurity staff and people collectively around the world.

- Linux Advisory Watch (<http://www.linuxsecurity.com/vuln-newsletter.html>) — A comprehensive newsletter that outlines the security vulnerabilities that have been announced throughout the week. It includes pointers to updated packages and descriptions of each vulnerability.
- Linux Security Week (<http://www.linuxsecurity.com/newsletter.html>) — The purpose of this document is to provide our readers with a quick summary of each week's most relevant Linux security headlines.
- Linux Security Discussion List (<http://www.linuxsecurity.com/general/maillinglists.html>) — This mailing list is for general security-related questions and comments.
- Linux Security Newsletters (<http://www.linuxsecurity.com/general/maillinglists.html>) — Subscription information for all newsletters.
- comp.os.linux.security FAQ (<http://www.linuxsecurity.com/docs/colsfaq.html>) — Frequently Asked Questions with answers for the comp.os.linux.security newsgroup.
- Linux Security Documentation (<http://www.linuxsecurity.com/docs/>) — A great starting point for information pertaining to Linux and Open Source security.

11.11.2. FTP Sites

CERT is the Computer Emergency Response Team. They often send out alerts of current attacks and fixes. See cert.org (<ftp://ftp.cert.org>) for more information.

ZEDZ (formerly Replay) (zedz.net (<http://www.zedz.net>)) has archives of many security programs. Since they are outside the US, they don't need to obey US crypto restrictions.

Matt Blaze is the author of CFS and a great security advocate. Matt's archive is available at att.com (<ftp://ftp.research.att.com/pub/mab>)

tue.nl is a great security FTP site in the Netherlands. [tue.nl](ftp://ftp.win.tue.nl/pub/security/) (<ftp://ftp.win.tue.nl/pub/security/>)

11.11.3. Web Sites

- The Hacker FAQ is a FAQ about hackers: The Hacker FAQ (<http://www.plethora.net/~seeb/faq/hacker.html>)
- The COAST archive has a large number of *UNIX* security programs and information: COAST (<http://www.cerias.purdue.edu/coast/>)
- SuSe Security Page: [suse security](http://www.suse.de/security/) (<http://www.suse.de/security/>)
- Rootshell.com is a great site for seeing what exploits are currently being used by crackers: rootshell (<http://www.rootshell.com/>)
- BUGTRAQ puts out advisories on security issues: BUGTRAQ archives (<http://online.securityfocus.com/archive/1>)
- CERT, the Computer Emergency Response Team, puts out advisories on common attacks on *UNIX* platforms: CERT home (<http://www.cert.org/>)
- Dan Farmer is the author of *SATAN* and many other security tools. His home site has some interesting security survey information, as well as security tools: <http://www.trouble.org/security> (<http://www.trouble.org/security>)
- The *GNU/Linux* security WWW is a good site for *GNU/Linux* security information: Linux Security WWW (<http://www.aoy.com/Linux/Security/>)
- Infilsec has a vulnerability engine that can tell you what vulnerabilities (<http://www.infilsec.com/vulnerabilities/>) affect a specific platform.
- CIAC sends out periodic security bulletins on common exploits: CIAC (<http://ciac.llnl.gov/cgi-bin/index/bulletins>)
- A good starting point for *GNU/Linux* Pluggable Authentication modules can be found at kernel.org (<http://www.kernel.org/pub/linux/libs/pam/>).
- WWW Security FAQ, written by Lincoln Stein, is a great web security reference. Find it at w3.org (<http://www.w3.org/Security/Faq/www-security-faq.html>)

11.11.4. Mailing Lists

Mandrake Linux security list: you can be informed for each security fix by subscribing to oursecurity mailing-list (<http://www.mandrakesecure.net/en/mlist.php>).

Bugtraq: To subscribe to bugtraq, send mail to listserv@netspace.org containing the message body "subscribe bugtraq". (See links above for archives.)

CIAC: Send e-mail to majordomo@tholia.llnl.gov. In the BODY (not subject) of the message put: "subscribe ciac-bulletin"

11.11.5. Books – Printed Reading Material

There are a number of good security books out there. This section lists a few of them. In addition to the security specific books, security is covered in a number of other books on system administration.

References

- D. Brent Chapman, Elizabeth D. Zwicky, *Building Internet Firewalls*, 1st Edition September 1995, ISBN 1-56592-124-0.
- Simson Garfinkel, Gene Spafford, *Practical UNIX & Internet Security*, 2nd Edition April 1996, ISBN 1-56592-148-8.
- Deborah Russell, G.T. Gangemi, Sr., *Computer Security Basics*, 1st Edition July 1991, ISBN 0-937175-71-4.
- Olaf Kirch, *Linux Network Administrator's Guide*, 1st Edition January 1995, ISBN 1-56592-087-2.
- Simson Garfinkel, *PGP: Pretty Good Privacy*, 1st Edition December 1994, ISBN 1-56592-098-8.
- David Icove, Karl Seger, William VonStorch, *Computer Crime A Crimefighter's Handbook*, 1st Edition August 1995, ISBN 1-56592-086-4.
- John S. Flowers, *Linux Security*, New Riders, March 1999, ISBN 0735700354.
- Anonymous, *Maximum Linux Security: A Hacker's Guide to Protecting Your Linux Server And Network*, July 1999, ISBN 0672313413.
- Terry Escamilla, *Intrusion Detection*, John Wiley and Sons, September 1998, ISBN 0471290009.
- Donn Parker, *Fighting Computer Crime*, John Wiley and Sons, September 1998, ISBN 0471163783.

11.12. Frequently Asked Questions

Q: Is it more secure to compile driver support directly into the kernel, instead of making it a module?

A: Some people think it is better to disable the ability to load device drivers using modules, because an intruder could load a Trojan module or a module that could affect system security.

However, in order to load modules, you must be root. The module object files are also only writable by root. This means the intruder would need root access to insert a module. If the intruder gains root access, there are more serious things to worry about than whether he will load a module.

Modules are for dynamically loading support for a particular device that may be infrequently used. On server computers, or firewalls for instance, this is very unlikely to happen. For this reason, it would make more sense to compile support directly into the kernel for machines acting as servers. Modules are also slower than support compiled directly in the kernel.

Q: Why does logging in as root from a remote machine always fail?

A: See *Root Security*, page 144. This is done intentionally to prevent remote users from attempting to connect via telnet to your computer as root, which is a serious security vulnerability, because then the root password would be transmitted, in clear text, across the network. Don't forget: potential intruders have time on their side, and can run automated programs to find your password. Additionally, this is done to keep a clear record of who logged in, not just root.

Q: How can I enable the *Apache* SSL extensions?

A: Simply install the package `mod_ssl`, and consult the documentation at `mod_ssl` home page (www.modssl.org).



You should also consider the `mod_sxnet` module, which is a plugin for `mod_ssl` and allows the activation of the “Thawte Secure Extranet”. `mod_ssl` encrypt communications, but `mod_ssl-sxnet` goes further and allows to securely authenticate the user of the web page thanks to a personal certificate. You have more info on this application on Thawte (<http://www.thawte.com/certs/strongextranet/>) or install the `mod_sxnet` module from your Mandrake distribution and read the included package documentation.

You might also try ZEDZ net (<http://www.zedz.net>) which has many pre-built packages, and is located outside of the United States.

Q: How can I manipulate user accounts, and still retain security?

A: Your **Mandrake Linux** distribution contains a great number of tools to change the properties of user accounts.

- The `pwconv` and `unpwconv` programs can be used to convert between shadowed and non-shadowed passwords.
- The `pwck` and `grpck` programs can be used to verify proper organization of the `/etc/passwd` and `/etc/group` files.
- The `useradd`, `usermod`, and `userdel` programs can be used to add, delete and modify user accounts. The `groupadd`, `groupmod`, and `groupdel` programs will do the same for groups.
- Group passwords can be created using `gpasswd`.

All these programs are “shadow-aware” – that is, if you enable shadow they will use `/etc/shadow` for password information, otherwise they won’t.

Q: How can I password-protect specific HTML documents using *Apache*?

A: I bet you didn’t know about <http://www.apacheweek.org> (<http://www.apacheweek.com>), did you?

You can find information on user authentication at [apacheweek](http://www.apacheweek.com/features/userauth) (<http://www.apacheweek.com/features/userauth>) as well as other web server-security tips from Apache (http://www.apache.org/docs/misc/security_tips.html).

11.13. Conclusion

By subscribing to the security alert mailing lists, and keeping current, you can do a lot towards securing your computer. If you pay attention to your log files and run something like `tripwire` regularly, you can do even more.

A reasonable level of computer security is not difficult to maintain on a home computer. More effort is required on business computers, but *GNU/Linux* can indeed be a secure platform. Due to the nature of *GNU/Linux* development, security fixes often come out much faster than they do on commercial operating systems, making *GNU/Linux* an ideal platform when security is a requirement.

Security-Related Terms

Included below are several of the most frequently used terms in computer security. A comprehensive dictionary of computer security terms is available in the LinuxSecurity.com Dictionary (<http://www.linuxsecurity.com/dictionary/>)

authentication

The process of knowing that the data received is the same as the data that was sent, and that the claimed sender is in fact the actual sender.

bastion Host

A computer system that must be highly secured because it is vulnerable to attack, usually because it is exposed to the Internet and is a main point of contact for users of internal networks. It gets its name from the highly fortified projects on the outer walls of medieval castles. Bastions overlook critical areas of defense, usually having strong walls, room for extra troops, and the occasional useful tub of boiling hot oil for discouraging attackers. Some reasonable definition here.

buffer overflow

Common coding style is to never allocate large enough buffers, and to not check for overflows. When such buffers overflow, the executing program (daemon or set-uid program) can be tricked in doing some other things. Generally this works by overwriting a function's return address on the stack to point to another location.

denial of service

An attack that consumes the resources on your computer for things it was not intended to be doing, thus preventing normal use of your network resources for legitimate purposes.

dual-homed Host

A general-purpose computer system that has at least two network interfaces.

firewall

A component or set of components that restricts access between a protected network and the Internet, or between other sets of networks.

host

A computer system attached to a network.

IP spoofing

IP Spoofing is a complex technical attack that is made up of several components. It is a security exploit that works by tricking computers in a trust relationship into thinking that you are someone that you really aren't. There is an extensive paper written by daemon9, route, and infinity in the Volume Seven, Issue Forty-Eight of Phrack Magazine.

non-repudiation

The property of a receiver being able to prove that the sender of some data did in fact send the data even though the sender might later deny ever having sent it.

packet

The fundamental unit of communication on the Internet.

packet filtering

The action a device takes to selectively control the flow of data to and from a network. Packet filters allow or block packets, usually while routing them from one network to another (most often from the Internet to an internal network, and vice-versa). To accomplish packet filtering, you set up rules that specify what types of packets (those to or from a particular IP address or port) are to be allowed and what types are to be blocked.

perimeter network

A network added between a protected network and an external network, in order to provide an additional layer of security. A perimeter network is sometimes called a DMZ.

proxy server

A program that deals with external servers on behalf of internal clients. Proxy clients talk to proxy servers, which relay approved client requests to real servers, and relay answers back to clients.

superuser

An informal name for root.

Chapter 12. Networking Overview

12.1. Copyright

This chapter is based on a *HOWTO* by Joshua D. Drake {POET}, which original is hosted by the The Linux Review (<http://www.thelinuxreview.com/>) web site.

The NET-3/4-HOWTO, NET-3, and Networking-HOWTO, information on how to install and configure networking support for Linux. Copyright (c) 1997 Terry Dawson, 1998 Alessandro Rubini, 1999 Joshua D. Drake {POET} - [thelinuxreview.com/](http://www.thelinuxreview.com/) (<http://www.thelinuxreview.com/>) is a FREE document. You may redistribute it under the terms of the GNU General Public License.

Modifications from v1.6.9, July 03, 2000, (C)opyright 2000 - 2002 MandrakeSoft.

12.2. How to Use this Chapter

This document is organized top-down. The first sections include informative material and can be skipped if you are not interested; what follows is a generic discussion of networking issues, and you must ensure you understand this before proceeding to more specific parts. The rest, technology-specific information, is grouped in three main sections: Ethernet and IP-related information, technologies pertaining to widespread PC hardware, and seldom-used technologies.

The suggested path through the document is thus the following:

Read the generic sections

These sections apply to every, or nearly every, technology described later. Thus, they are very important for you to understand. On the other hand, I expect many of the readers to be already confident with this material.

Consider your network

You should know how your network is, or will be, designed and exactly what hardware and technology types you will be implementing.

Read the *Ethernet Information*, page 179 if you are directly connected to a LAN or the Internet:

This section describes basic Ethernet configuration and the various features that Linux offers for IP networks, like firewalling, advanced routing and so on.

Read the next section if you are interested in low-cost local networks or dial-up connections

The section describes PLIP, PPP, SLIP and ISDN, the widespread technologies used on personal workstations.

Read the technology-specific sections related to your requirements

If your needs differ from IP and/or common hardware, the final section covers details specific to non-IP protocols and peculiar communication hardware.

Do the configuration work

You should actually try to configure your network and take careful note of any problems you have.

Look for further help if needed

If you experience problems that this document does not help you to resolve, then read the section related to where to get help or where to report bugs.

Have fun!

Networking is fun, enjoy it.

12.2.1. Conventions Used in this Document

No special convention is used here, but you must be warned about the way commands are shown. Following the classic *UNIX* documentation, any command you should type to your shell is prefixed by a prompt. This howto shows “user%” as the prompt for commands that do not require superuser privileges, and “root#” as the prompt for commands that need to be run as root. I chose to use “root#” instead of a plain “#” to prevent confusion with snapshots from shell scripts, where the hash mark is used to define comment lines.

When “Kernel Compile Options” are shown, they are represented in the format used by **menuconfig**. They should be understandable even if you (like me) are not used to **menuconfig**. If you are in doubt about the options’ nesting, running the program once can’t do anything, but help.

12.3. General Information about Linux Networking

12.3.1. Linux Networking Resources

There are a number of places where you can find good information about Linux networking.

There is a wealth of consultants available. A searchable listing can be found on the thelinuxreview.com/ (<http://www.thelinuxreview.com/>) web site.

Alan Cox, the current maintainer of the Linux kernel networking code, maintains a world wide web page that contains highlights of current and new developments in Linux Networking at: [www.linux.org.uk](http://www.linux.org.uk/NetNews.html) (<http://www.linux.org.uk/NetNews.html>).

There is a newsgroup in the Linux news hierarchy dedicated to networking and related matters, it is: comp.os.linux.networking (news:comp.os.linux.networking).

You can also subscribe to a mailing list where you may ask questions relating to Linux networking. To subscribe, you should send a mail message:

```
To: majordomo@vger.rutgers.edu
Subject: anything at all
Message:
subscribe linux-net
```

Please remember when reporting any problem to include as much relevant detail about the problem as you can. Specifically, you should identify the versions of software that you are using, especially the kernel version, the version of tools such as pppd/ or dip and the exact nature of the problem you are experiencing. This means taking note of the exact syntax of any error message you receive and of any command you are issuing.

12.3.2. Where to Get some non Linux-Specific Network Information

If you are after some basic tutorial information on TCP/IP networking generally, then I recommend you take a look at the following documents:

TCP/IP Introduction

This document comes as both a text (<ftp://athos.rutgers.edu/runet/tcp-ip-intro.doc>) and a postscript (<ftp://athos.rutgers.edu/runet/tcp-ip-intro.ps>) version.

TCP/IP Administration

This document comes both as a text (<ftp://athos.rutgers.edu/runet/tcp-ip-admin.doc>) and a postscript (<ftp://athos.rutgers.edu/runet/tcp-ip-admin.ps>) version.

If you are looking for some more detailed information on TCP/IP networking, then I highly recommend:

“ **Internet working with TCP/IP, Volume 1: principles, protocols and architecture**, by Douglas E. Comer, ISBN 0-13-227836-7, Prentice Hall publications, Third Edition, 1995.”

If you want to learn about how to write network applications in a *UNIX*-compatible environment, then I also highly recommend:

“ **Unix Network Programming**, by W. Richard Stevens, ISBN 0-13-949876-1, Prentice Hall publications, 1990.”

A second edition of this book is appearing on the bookshelves; the new book is made up of three volumes: check Prentice-Hall's (<http://www.phptr.com/>) web site for more information.

You might also try the `comp.protocols.tcp-ip` (`news:comp.protocols.tcp-ip`) newsgroup.

An important source of specific technical information relating to the Internet and the TCP/IP suite of protocols are RFCs. RFC is an acronym for “Request For Comment” and is the standard means of submitting and documenting Internet protocol standards. There are many RFC repositories. Many of these sites use the ftp protocol and others provide World Wide Web access with an associated search engine that allows you to search the RFC database for particular keywords.

One possible source for RFCs is on the Nexor RFC database (<http://pubweb.nexor.co.uk/public/rfc/index/rfc.html>) web site.

12.4. Generic Network Configuration Information

You will need to know and understand the following subsections before you actually try to configure your network. They are fundamental principles that apply regardless of the exact nature of the network you wish to deploy.

12.4.1. What Do I Need to Start?

Before you start building or configuring your network, you will need some things. The most important of these are:

12.4.1.1. Current Kernel Source (Optional)



Your **Mandrake Linux** distribution comes with networking enabled, therefore it may not be required to recompile the kernel. If you are running well known hardware, you should be just fine. For example: 3COM NIC, NE2000 NIC, or an Intel NIC. However, if you find yourself in the position that you do need to update the kernel, the following information is provided.

Because the kernel you are running now might not yet have support for the network types or cards that you wish to use, you will probably need the kernel source in order to recompile the kernel with the appropriate options.

However, as long as you stay within the mainstream of hardware, there should be no need to recompile your kernel unless there is a very specific feature that you need.

You can always obtain the latest kernel source from the `sunsite.unc.edu` (<ftp://sunsite.unc.edu/pub/linux/kernel.org/pub/linux/kernel/>) web site. This is not the official site, but they have LOTS of bandwidth and capacity. The official site is `kernel.org` but please use the above if you can. Please remember that `ftp.kernel.org` is seriously overloaded. Use a mirror instead.

Normally, the kernel source will be untarred into the `/usr/src/linux` directory. For information on how to apply patches and build the kernel, you should read the Kernel-HOWTO (<http://linuxdoc.org/HOWTO/Kernel-HOWTO.html>). For information on how to configure kernel modules, you should read the *Modules mini-HOWTO*. Also, the README file found in the kernel sources and the Documentation directory are very informative for the brave reader.

Unless specifically stated otherwise, I recommend you stick with the standard kernel release (the one with the even number as the second digit in the version number). Development release kernels (the ones with the odd second digit) may have structural or other changes that may cause problems working with the other software on your system. If you are uncertain that you could resolve those sorts of problems, in addition to the potential for there being other software errors, then don't use them.

12.4.1.2. IP Addresses, an Explanation

Internet Protocol (IP) addresses are composed of four bytes¹. The convention is to write addresses in what is called “dotted decimal notation”. In this form, each byte is converted to a decimal number, (0-255), dropping any leading zeroes, unless the number is zero and written with each byte separated by a “.” character. By convention, each host or router interface has an IP address. It is legal for the same IP address to be used on each interface of a single machine in some circumstances, but usually, each interface will have its own address.

Internet Protocol networks are contiguous sequences of IP addresses. All addresses within a network have a number of digits within the address in common. The portion of the address that is common amongst all addresses within the network is called the “network portion” of the address. The remaining digits are called the “host portion”. The number of bits that are shared by all addresses within a network is called the netmask and it is the latter’s role to determine which addresses belong to the network it is applied to and which don’t. For example, consider the following:

Host Address	192.168.110.23
Network Mask	255.255.255.0
Network Portion	192.168.110.
Host Portion	.23
Network Address	192.168.110.0
Broadcast Address	192.168.110.255

Any address that is “bitwise anded” with its netmask will reveal the address of the network it belongs to. The network address is therefore always the lowest numbered address within the range of addresses on the network and always has the host portion of the address coded all zeroes.

The broadcast address is a special one which every host on the network listens to, in addition to its own unique address. This address is the one that datagrams are sent to if every host on the network is meant to receive it. Certain types of data like routing information and warning messages are transmitted to the broadcast address so that every host on the network can receive it simultaneously. There are two commonly used standards for what the broadcast address should be. The most widely accepted one is to use the highest possible address on the network as the broadcast address. In the example above, this would be 192.168.110.255. For some reason, other sites have adopted the convention of using the network address as the broadcast address. In practice, it doesn’t matter very much which you use, but you must make sure that every host on the network is configured with the same broadcast address.

For administrative reasons, some time early in the development of the IP protocol, some arbitrary groups of addresses were formed into networks and these networks were grouped into what are called classes. These classes provide a number of standard size networks that could be allocated. The ranges allocated are:

Network Class	Netmask	Network Addresses
A	255.0.0.0	0.0.0.0 - 127.255.255.255
B	255.255.0.0	128.0.0.0 - 191.255.255.255
C	255.255.255.0	192.0.0.0 - 223.255.255.255
Multicast	240.0.0.0	224.0.0.0 - 239.255.255.255

What addresses you should use depends on exactly what it is that you are doing. You may have to use a combination of the following activities to get all the addresses you need:

Installing a Linux machine on an existing IP network

If you wish to install a Linux machine onto an existing IP network, then you should contact whoever administers the network and ask them for the following information:

- Host IP Address
- IP Network Address

1. For the version 4 of IP, a.k.a IPv4

- IP Broadcast Address
- IP Netmask
- Router Address
- Domain Name Server Address

You should then configure your Linux network device with those details. You can not make them up and expect your configuration to work.

Building a brand new network that will never connect to the Internet

If you are building a private network and you never intend that network to be connected to the Internet, then you can choose whatever addresses you like. However, for safety and consistency reasons, there have been some IP network addresses that have been reserved specifically for this purpose. These are specified in RFC1597 and are as follows:

Network Class	Netmask	Network Addresses
A	255.0.0.0	10.0.0.0 - 10.255.255.255
B	255.255.0.0	172.16.0.0 - 172.31.255.255
C	255.255.255.0	192.168.0.0 - 192.168.255.255

Table 12-1. Reserved Private Network Allocations

You should first decide how large you want your network to be, and then choose as many addresses as you require.

12.4.2. Routing

Routing is a big topic. It is easily possible to write large volumes of text about it. Most of you will have fairly simple routing requirements, some of you will not. I will cover some basic fundamentals of routing only. If you are interested in more detailed information, then I suggest you refer to the references provided at the start of the document.

Let's start with a definition. What is IP routing? Here is one that I'm using:

"IP routing is the process by which a host with multiple network connections decides where to deliver IP datagrams it has received."

It might be useful to illustrate this with an example. Imagine a typical office router, it might have a PPP link off the Internet, a number of Ethernet segments feeding the workstations, and another PPP link off to another office. When the router receives a datagram on any of its network connections, routing is the mechanism that it uses to determine which interface it should send the datagram to next. Simple hosts also need to route, all Internet hosts have two network devices, one is the loopback interface described above, and the other is the one it uses to talk to the rest of the network, perhaps an Ethernet, perhaps a PPP or SLIP serial interface.

Ok, so how does routing work? Each host keeps a special list of routing rules, called a routing table. This table contains rows which typically contain at least three fields: the first is a destination address, the second is the name of the interface to which the datagram is to be routed, and the third is optionally the IP address of another machine which will carry the datagram on its next step through the network. With Linux, you can see this table by using the following command:

```
user% cat /proc/net/route
```

or by using either one of the following commands:

```
user% /sbin/route -n
user% netstat -r
```

The routing process is fairly simple: an incoming datagram is received, the destination address (who it is for) is examined and compared with each entry in the table. The entry that best matches that address is selected and the datagram is forwarded to the specified interface. If the gateway field is filled, then the datagram is forwarded to that host via the specified interface. Otherwise, the destination address is assumed to be on the network supported by the interface.

12.4.2.1. What Does the Routed Program Do?

The routing configuration described above is best suited for simple network arrangements where there is only one possible path to a determined destination. When you have a more complex network arrangement, things get a little more complicated. Fortunately for most of you, this won't be an issue.

The big problem with "manual routing" or "static routing" as described, is that if a machine or link fails in your network, then the only way you can direct your datagrams another way, if another way exists, is by manually intervening and executing the appropriate commands. Naturally this is clumsy, slow, impractical and hazard prone. Various techniques have been developed to automatically adjust routing tables in the event of network failures where there are alternate routes. All of these techniques are loosely grouped by the term "dynamic routing protocols".

You may have heard of some of the more common dynamic routing protocols. The most common are probably RIP (Routing Information Protocol) and OSPF (Open Shortest Path First Protocol). The Routing Information Protocol is very common on small networks such as small-to-medium size corporate networks or building networks. OSPF is more modern and more capable of handling large network configurations and better suited to environments where there is a large number of possible paths through the network. Common implementations of these protocols are: **routed** - RIP and **gated** - RIP, OSPF and others. The routed program is normally supplied with your Linux distribution or is included in the "NetKit" package detailed above.

An example of where and how you might use a dynamic routing protocol might look something like figure 12-1.




Figure 12-1. A Dynamic Routing Example

We have three routers: A, B and C. Each one supports one Ethernet segment with a Class C IP network (net-mask 255.255.255.0). Each router also has a PPP link to each of the other routers. The network forms a triangle.

It should be clear that the routing table at router A could look like:

```

root# route add -net 192.168.1.0 netmask 255.255.255.0 eth0
root# route add -net 192.168.2.0 netmask 255.255.255.0 ppp0
root# route add -net 192.168.3.0 netmask 255.255.255.0 ppp1

```

This would work just fine until the link between router A and B should fail. If that link fails, then with the routing entry shown above, hosts on the Ethernet segment of A could not reach hosts on the Ethernet segment on B because their datagram would be directed to router A's ppp0 link, which is broken. They could still continue to talk to hosts on the Ethernet segment of C and hosts on the C's Ethernet segment could still talk to hosts on B's Ethernet segment, because the link between B and C is still intact.

But wait, if A can talk to C and C can still talk to B, why shouldn't A route its datagrams for B via C and let C send them to B? This is exactly the sort of problem that dynamic routing protocols like RIP were designed to solve. If each of the routers A, B and C were running a routing daemon, then their routing tables would be automatically adjusted to reflect the new state of the network should any one of the links in the network fail. To configure such a network is simple. For each router, you only need to do two things. In this case, for router A:

```

root# route add -net 192.168.1.0 netmask 255.255.255.0 eth0
root# /usr/sbin/routed

```

The **routed** routing daemon automatically finds all active network ports when it starts and sends and listens for messages on each of the network devices, to allow it to determine and update the routing table on the host.

This has been a very brief explanation of dynamic routing and where you would use it. If you want more information, then you should refer to the suggested references listed at the top of the document.

The important points relating to dynamic routing are:

1. You only need to run a dynamic routing protocol daemon when your Linux machine has the possibility of selecting multiple possible routes to a destination. An example of this would be if you plan to use IP masquerading.
2. The dynamic routing daemon will automatically modify your routing table to adjust to changes in your network.
3. RIP is suited for small-to-medium size networks.

12.5. Ethernet Information

This section covers information specific to Ethernet and the configuring of Ethernet cards.

12.5.1. Supported Ethernet Cards

12.5.1.1. 3Com

- 3Com 3c501 - (avoid like the plague) (3c501 driver);
- 3Com 3c503 (3c503 driver), 3c505 (3c505 driver), 3c507 (3c507 driver), 3c509/3c509B (ISA) / 3c579 (EISA);
- 3Com Etherlink III Vortex Ethercards (3c590, 3c592, 3c595, 3c597) (PCI), 3Com Etherlink XL Boomerang (3c900, 3c905) (PCI) and Cyclone (3c905B, 3c980) Ethercards (3c59x driver) and 3Com Fast EtherLink Ethercard (3c515) (ISA) (3c515 driver);
- 3Com 3ccfe575 Cyclone Cardbus (3c59x driver);
- 3Com 3c575 series Cardbus (3c59x driver) (most PCMCIA cards should be detected);

12.5.1.2. AMD, ATT, Allied Telesis, Ansel, Apricot

- AMD LANCE (79C960) / PCnet-ISA/PCI (AT1500, HP J2405A, NE1500/NE2100);
- ATT GIS WaveLAN;
- Allied Telesis AT1700;
- Allied Telesis LA100PCI-T;
- Allied Telesyn AT2400T/BT (“ne” module);
- Ansel Communications AC3200 (EISA);
- Apricot Xen-II / 82596.

12.5.1.3. Cabletron, Cogent, Crystal LAN

- Cabletron E21xx;
- Cogent EM110;
- Crystal LAN CS8920, Cs8900.

12.5.1.4. Danpex, DEC, Digi, DLink

- Danpex EN-9400;
- DEC DE425 (EISA) / DE434/DE435 (PCI) / DE450/DE500 (DE4x5 driver);
- DEC DE450/DE500-XA (dc21x4x) (Tulip driver);
- DEC DEPCA and EtherWORKS;
- DEC EtherWORKS 3 (DE203, DE204, DE205);
- DECchip DC21x4x “Tulip”;
- DEC QSilver’s (Tulip driver);
- Digi International RightSwitch;
- DLink DE-220P, DE-528CT, DE-530+, DFE-500TX, DFE-530TX.

12.5.1.5. Fujitsu, HP, ICL, Intel

- Fujitsu FMV-181/182/183/184;
- HP PCLAN (27245 and 27xxx series);
- HP PCLAN PLUS (27247B and 27252A);
- HP 10/100VG PCLAN (J2577, J2573, 27248B, J2585) (ISA/EISA/PCI);
- ICL EtherTeam 16i / 32 (EISA);
- Intel EtherExpress;
- Intel EtherExpress Pro.

12.5.1.6. KTI, Macromate, NCR NE2000/1000, Netgear, New Media

- KTI ET16/P-D2, ET16/P-DC ISA (work jumperless and with hardware-configuration options);
- Macromate MN-220P (PnP or NE2000 mode);
- NCR WaveLAN;
- NE2000/NE1000 (be careful with clones);
- Netgear FA-310TX (Tulip chip);
- New Media Ethernet.

12.5.1.7. PureData, SEEQ, SMC

- PureData PDUC8028, PDI8023;
- SEEQ 8005;
- SMC Ultra / EtherEZ (ISA);
- SMC 9000 series;
- SMC PCI EtherPower 10/100 (DEC Tulip driver);
- SMC EtherPower II (epic100.c driver).

12.5.1.8. Sun Lance, Sun Intel, Schneider, WD, Zenith, IBM, Enyx

- Sun LANCE adapters (kernel 2.2 and newer);
- Sun Intel adapters (kernel 2.2 and newer);
- Schneider and Koch G16;
- Western Digital WD80x3;
- Zenith Z-Note / IBM ThinkPad 300 built-in adapter;
- Znyx 312 etherarray (Tulip driver);

12.5.2. General Ethernet Information

Ethernet device names are `eth0`, `eth1`, `eth2` etc. The first card detected by the kernel is assigned `eth0` and the rest are assigned sequentially in the order they are detected.

Once you have your kernel properly built to support your Ethernet card, the card configuration is easy.

Typically, you would use something like (which most distributions already do for you, if you configured them to support your Ethernet):

```
root# ifconfig eth0 192.168.0.1 netmask 255.255.255.0 up
root# route add -net 192.168.0.0 netmask 255.255.255.0 eth0
```

Most of the Ethernet drivers were developed by Donald Becker (<mailto:becker@CESDIS.gsfc.nasa.gov>).

12.5.3. Using 2 or More Ethernet Cards in The Same Machine

The module will typically detect all of the installed cards.

Detection information is stored in the `/etc/conf.modules` file.

Consider that a user has 3 NE2000 cards, one at 0x300, one at 0x240, and one at 0x220. You would add the following lines to the `/etc/conf.modules` file:

```
alias eth0 ne
alias eth1 ne
alias eth2 ne
options ne io=0x220,0x240,0x300
```

What this does is tell the **modprobe** program to look for 3 NE based cards at the following addresses. It also states in which order they should be found and the device they should be assigned.

Most ISA modules can take multiple comma separated I/O values. For example:

```
alias eth0 3c501
alias eth1 3c501
options eth0 -o 3c501-0 io=0x280 irq=5
options eth1 -o 3c501-1 io=0x300 irq=7
```

The `-o` option allows for a unique name to be assigned to each module. The reason for this is that you can not load two copies of the same module.

The `irq=` option is used to specify the hardware IRQ, and the `io=` to specify the different io ports.

By default, the Linux kernel only probes for one Ethernet device. You need to pass command-line arguments to the kernel in order to force detection of further boards.

To learn how to make your Ethernet card(s) working under Linux, you should refer to the Ethernet-HOWTO (<http://linuxdoc.org/HOWTO/Ethernet-HOWTO.html>).

12.6. IP-Related Information

These sections cover information specific to IP.

12.6.1. DNS

DNS stands for Domain Name System. It is the system responsible for mapping a machine name such as `www.mandrakesoft.com` with the IP address of that machine, in this case: 216.71.116.162 at the time of writing. With DNS, mapping is available in both directions, that is from name to IP and vice versa.

The DNS is composed of a great number of machines all over the Internet responsible for a certain number of names. Each machine is attributed a DNS server to which it can ask to map a particular name with its address. If that server does not have the answer, then it asks to another one and so on. You can also have a local DNS responsible for mapping addresses on your LAN.

We can differentiate two major DNS classes: caching DNS and master DNS servers. The first one only “remembers” a previous request and then can answer it without asking a master DNS server once more. The latter servers are really responsible as a last resort to map an address with a name - or possibly specify that a given name does not map any address.

12.6.2. DHCP And DHCPD

DHCP is an acronym for Dynamic Host Configuration Protocol. The creation of DHCP has made configuring the network on multiple hosts extremely simple. Instead of having to configure each host separately, you can assign all of the commonly used parameters by the hosts using a DHCP server.

Each time the host boots up, it will broadcast a packet to the network. This packet is a call to any DHCP servers located on the same segment to configure the host.

DHCP is extremely useful in assigning items such as the IP address, netmask, and gateway of each host.

12.7. Using Common PC Hardware

12.7.1. ISDN

The Integrated Services Digital Network (ISDN) is a series of standards that specify a general purpose-switched digital data network. An ISDN “call” creates a synchronous point-to-point data service to the destination. ISDN is generally delivered on a high-speed link that is broken down into a number of discrete channels. There are two different types of channels: the “B Channels” which will actually carry the user data, and a single channel called the “D channel” which is used to send control information to the ISDN exchange to establish calls and other functions. In Australia for example, ISDN may be delivered on a 2Mbps link that is broken into 30 discrete 64kbps B channels with one 64kbps D channel. Any number of channels may be used at a time and in any combination. You could for example establish 30 separate calls to 30 different destinations at 64kbps each, or you could establish 15 calls to 15 different destinations at 128kbps each (two channels used per call), or just a small number of calls and leave the rest idle. A channel may be used for either incoming or outgoing calls. The original intention of ISDN was to allow telecommunications companies to provide a single data service which could deliver either telephone (via digitized voice) or data services to your home or business without requiring you to make any special configuration changes.

There are a few different ways to connect your computer to an ISDN service. One way is to use a device called a “Terminal Adaptor” which plugs into the Network Terminating Unit that your telecommunications carrier will have installed when you got your ISDN service and presents a number of serial interfaces. One of the latter is used to enter commands to establish calls and configuration while the others are actually connected to the network devices that will use the data circuits when they are established. Linux will work in this sort of configuration without modification, you just treat the port on the Terminal Adaptor like you would treat any other serial device. Another way, which is the way the kernel ISDN support is designed for, allows you to install an ISDN card into your Linux machine and then has your Linux software handle the protocols and make the calls itself.

Kernel Compile Options:

```
ISDN subsystem --->
<*> ISDN support
[ ] Support synchronous PPP
[ ] Support audio via ISDN
< > ICN 2B and 4B support
< > PCBIT-D support
< > Teles/NICCY1016PC/Creatix support
```

The Linux implementation of ISDN supports a number of different types of internal ISDN cards. Hereunder are those listed in the kernel configuration options:

- ICN 2B and 4B;
- Octal PCBIT-D;
- Teles ISDN-cards and compatibles.

Some of these cards require software to be downloaded in order to make them operational. There is a separate utility to do this with.

Full details on how to configure the Linux ISDN support is available from the `/usr/src/linux/Documentation/isdn/` directory and a FAQ dedicated to **isdn4linux** is available on the [www.lrz-muenchen.de](http://www.lrz-muenchen.de/~ui161ab/www/isdn/) (<http://www.lrz-muenchen.de/~ui161ab/www/isdn/>) web site. (You can click on the English flag to get an English version).



About PPP. The PPP suite of protocols will operate over both asynchronous or synchronous serial lines. The commonly distributed PPP daemon for Linux, `pppd`, supports only asynchronous mode. If you wish to run the PPP protocol over your ISDN service, you need a specially modified version. Details of where to find it are available in the documentation referred to above.

12.7.2. PLIP

During development of the 2.1 kernel versions, support for the parallel port was changed to a better setup.

Kernel Compile Options:

```
General setup --->
  [*] Parallel port support
Network device support --->
  <*> PLIP (parallel port) support
```

The new code for PLIP behaves like the old one. Use the same **ifconfig** and **route** commands as in the previous section, but initialization of the device is different due to the advanced parallel port support.

The “first” PLIP device is always called `plip0`, where first is the first device detected by the system, similarly to what happens for Ethernet devices. The actual parallel port being used is one of the available ports, as shown in `/proc/parport`. For example, if you have only one parallel port, you’ll only have a directory called `/proc/parport/0`.

If your kernel didn’t detect the IRQ number used by your port, `insmod plip` will fail. In this case, just write the right number to `/proc/parport/0/irq` and reinvoke **insmod**.

Complete information about parallel port management is available in the `Documentation/parport.txt` file, part of your kernel sources.

12.7.3. PPP

Due to the nature of PPP, its size, complexity, and flexibility, it has been moved to its own *HOWTO*. The PPP-HOWTO is still a Linux Documentation Project document (<http://www.linuxdoc.org>), but its official home is now on the [thelinuxreview.com](http://www.thelinuxreview.com) (<http://www.thelinuxreview.com>) web site, in the PPP section (<http://www.thelinuxreview.com/howto/ppp>).

12.8. Other Network Technologies

The following subsections are specific to particular network technologies. The information contained in these sections does not necessarily apply to any other type of network technology. The topics are sorted alphabetically.

12.8.1. ARCNet

ARCNet device names are `arc0e`, `arc1e`, `arc2e` etc. or `arc0s`, `arc1s`, `arc2s` etc. The first card detected by the kernel is assigned `arc0e` or `arc0s` and the rest are assigned sequentially in the order they are detected. The letter at the end signifies whether you selected the Ethernet encapsulation packet format or the RFC1051 packet format.

Kernel Compile Options:

```
Network device support --->
  [*] Network device support
    <*> ARCnet support
      [ ] Enable arc0e (ARCnet "Ether-Encap" packet format)
      [ ] Enable arc0s (ARCnet RFC1051 packet format)
```

Once you have your kernel properly built to support your Ethernet card, then configuring your card is easy.

Typically, you would use something like:

```
root# ifconfig arc0e 192.168.0.1 netmask 255.255.255.0 up
root# route add -net 192.168.0.0 netmask 255.255.255.0 arc0e
```

Please refer to the `/usr/src/linux/Documentation/networking/arcnet.txt` and `/usr/src/linux/Documentation/networking/arcnet-hardware.txt` files for further information.

ARCNet support was developed by Avery Pennarun, `apenwarr@foxnet.net`.

12.8.2. Appletalk (AF_APPLETALK)

The Appletalk support has no special device names as it uses existing network devices.

Kernel Compile Options:

```
Networking options --->
  <*> Appletalk DDP
```

Appletalk support allows your Linux machine to interwork with Apple networks. An important use for this is to share resources such as printers and disks between both your Linux and Apple computers. Additional software is required, this is called **netatalk**. Wesley Craig `netatalk@umich.edu` represents a team called the “Research Systems Unix Group” at the University of Michigan and they have produced the **netatalk** package which provides software that implements the Appletalk protocol stack and some useful utilities. The **netatalk** package will either have been supplied with your Linux distribution, or you will have to ftp it from its home site at the University of Michigan (`ftp://terminator.rs.itd.umich.edu/unix/netatalk/`).

To build and install the package, you should do:

```
user% tar xvfz ../netatalk-1.4b2.tar.Z
user% make
root# make install
```

You may want to edit the Makefile before calling `make` to actually compile the software. Specifically, you might want to change the `DESTDIR` variable which defines where the files will be installed later. The default of `/usr/local/atalk` is fairly safe.

12.8.2.1. Configuring the Appletalk Software

The first thing you need to do to make it all work is to ensure that the appropriate entries in the `/etc/services` file are present. The entries you need are:

```
rtmp 1/ddp # Routing Table Maintenance Protocol
nbp 2/ddp # Name Binding Protocol
echo 4/ddp # AppleTalk Echo Protocol
zip 6/ddp # Zone Information Protocol
```

The next step is to create the Appletalk configuration files in the `/usr/local/atalk/etc` directory (or wherever you installed the package).

The first file to create is the `/usr/local/atalk/etc/atalkd.conf` file. Initially, this file needs only one line that gives the name of the network device that supports the network that your Apple machines are on:

```
eth0
```

The Appletalk daemon program will add extra details after it is run.

12.8.2.2. Exporting a Linux Filesystem Via Appletalk

You can export filesystems from your Linux machine to the network so that Apple machines on the network can share them.

To do this you must configure the `/usr/local/ataalk/etc/AppleVolumes.system` file. There is another configuration file called `/usr/local/ataalk/etc/AppleVolumes.default` which has exactly the same format and describes which filesystems users connecting with guest privileges will receive.

Full details on how to configure these files and what the various options are can be found in the `afpd` man page: `afpd`.

A simple example might look like...

```
/tmp Scratch
/home/ftp/pub "Public Area"
```

... which would export your `/tmp` filesystem as AppleShare Volume “Scratch” and your ftp public directory as AppleShare Volume “Public Area”. The volume names are not mandatory, the daemon will choose some for you, but it won’t hurt to specify them anyway.

12.8.2.3. Sharing your Linux Printer across Appletalk

You can share your Linux printer with your Apple machines quite simply. You need to run the **papd** program which is the Appletalk Printer Access Protocol Daemon. When you run this program, it will accept requests from your Apple machines and spool the print jobs to your local line printer daemon for printing.

You need to edit the `/usr/local/ataalk/etc/papd.conf` file to configure the daemon. The syntax of this file is the same as that of your usual `/etc/printcap` file. The name you give to the definition is registered with the Appletalk naming protocol, NBP.

A sample configuration might look like...

```
TricWriter:\
:pr=lp:op=cg:
```

... which would make a printer named “TricWriter” available to your Appletalk network, and all accepted jobs would be printed to the lp Linux printer (as defined in the `/etc/printcap` file), using *CUPS*. The entry `op=cg` says that the Linux user `cg` is the printer operator.

12.8.2.4. Starting the AppleTalk Software

Ok, you should now be ready to test this basic configuration. There is an `rc.ataalk` file supplied with the `netatalk` package that should work well for you, so all you should have to do is:

```
root# /usr/local/ataalk/etc/rc.ataalk
```

... and all should start up and run correctly. You should see no error messages and the software will send messages to the console indicating each stage as it starts.

12.8.2.5. Testing the AppleTalk Software

To test that the software is functioning properly, go to one of your Apple machines, pull down the Apple menu, select the Chooser, click on AppleShare, and your Linux box should appear.

12.8.2.6. Caveats of the AppleTalk Software

1. You may need to start the Appletalk support before you configure your IP network. If you have problems starting the Appletalk programs, or if after you start them you have trouble with your IP network, then try starting the Appletalk software before you run your `/etc/rc.d/rc.inet1` file.
2. The **afpd** (Apple Filing Protocol Daemon) severely messes up your hard disk. Below the mount points, it creates a couple of directories called `.AppleDesktop` and `NetworkTrashFolder`. Then, for each directory you access, it will create a `.AppleDouble` below it so it can store resource forks, etc. So think twice before exporting `/`, you will have a great time cleaning up afterwards.
3. The **afpd** program expects clear text passwords from the Macs. Security could be a problem, so be very careful when you run this daemon on a machine connected to the Internet. You have yourself to blame if somebody nasty does something bad.
4. The existing diagnostic tools such as **netstat** and **ifconfig** don't support Appletalk. The raw information is available in the `/proc/net/` directory if you need it.

12.8.2.7. More Information

For a much more detailed description of how to configure Appletalk for Linux, refer to Anders Brownworth *Linux Netatalk-HOWTO* page on the TheHamptons.com (<http://thehamptons.com/anders/netatalk/>) web site.

12.8.3. ATM

Werner Almesberger (<werner.almesberger@lrc.di.epfl.ch>) manages a project which goal is to provide Asynchronous Transfer Mode support for Linux. Current information on the status of the project may be obtained the ATM on Linux (<http://linux-atm.sourceforge.net/>) web site.

12.8.4. AX25 (AF_AX25)

AX.25 device names are `sl0`, `sl1`, etc. in 2.0.* kernels or `ax0`, `'ax1`, etc. in 2.1.* kernels.

Kernel Compile Options:

```
Networking options --->
  [*] Amateur Radio AX.25 Level 2
```

The AX25, Netrom and Rose protocols are covered by the AX25-HOWTO (<http://linuxdoc.org/HOWTO/AX25-HOWTO.html>). These protocols are used by Amateur Radio Operators world wide in packet radio experimentation.

Most of the work for implementation of these protocols has been done by Jonathon Naylor, jsn@cs.nott.ac.uk, and the new *HOWTO* maintainer is Jeff Tranter (tranter@pobox.com).

12.8.5. DECNet

Support for DECNet is now included in current stable kernel (2.4). With **Mandrake Linux**, it has been available since kernel 2.2.

12.8.6. FDDI

FDDI device names are `fddi0`, `fddi1`, `fddi2` etc. The first card detected by the kernel is assigned `fddi0` and the rest are assigned sequentially in the order they are detected.

Larry Stefani (l1stefani@ultranet.com) has developed a driver for the Digital Equipment Corporation FDDI EISA and PCI cards.

Kernel Compile Options:

```
Network device support --->
  [*] FDDI driver support
  [*] Digital DEFEA and DEFPA adapter support
```

If your kernel is built to support the FDDI driver and installed, configuration of the FDDI interface is almost identical to that of an Ethernet interface. You just specify the appropriate FDDI interface name in the **ifconfig** and **route** commands.

12.8.7. Frame Relay

The Frame Relay device names are `dlci00`, `dlci01` etc., for the DLCI encapsulation devices and `sdla0`, `sdla1` etc., for the FRAD(s).

Frame relay is a new networking technology that is designed to suit data communications traffic that is of a “bursty” or intermittent nature. You connect to a frame relay network using a Frame Relay Access Device (FRAD). The Linux frame relay supports IP over frame relay as described in RFC-1490.

Kernel Compile Options:

```
Network device support --->
  <*> Frame relay DLCI support (EXPERIMENTAL)
  (24)  Max open DLCI
  (8)   Max DLCI per device
  <*>  SDLA (Sangoma S502/S508) support
```

Mike McLagan (mike.mclagan@linux.org) developed the frame relay support and configuration tools.

Currently, the only FRAD I know of that are supported are the Sangoma Technologies (<http://www.sangoma.com/>) S502A, S502E and S508, as well as the one from Emerging Technologies (<http://www.etinc.com/>).

To configure the FRAD and DLCI devices after you have rebuilt your kernel, you will need the frame relay configuration tools. These are available from the [ftp.invlogic.com](ftp://ftp.invlogic.com/pub/linux/fr/) (<ftp://ftp.invlogic.com/pub/linux/fr/>) site.

Compiling and installing the tools is straightforward, but the lack of a top level Makefile makes it a fairly manual process:

```
user% tar xvfz ../frad-0.15.tgz
user% cd frad-0.15
user% for i in common dlci frad; do make -C $i clean; make -C $i; done
root# mkdir /etc/frad
root# install -m 644 -o root -g root bin/*.sfm /etc/frad
root# install -m 700 -o root -g root frad/fradcfg /sbin
root# install -m 700 -o root -g root dlci/dlcicfg /sbin
```

Note that the previous commands use **sh** syntax. If you use a **cs**h flavor instead (like **tcsh**), the **for** loop will look different.

After installing the tools, you need to create an `/etc/frad/router.conf` file. You can use this template, which is a modified version of one of the example files:

```
# /etc/frad/router.conf
# This is a template configuration for frame relay.
```

```

# All tags are included. The default values are based on the code
# supplied with the DOS drivers for the Sangoma S502A card.
#
# A '#' anywhere in a line constitutes a comment
# Blanks are ignored (you can indent with tabs too)
# Unknown [] entries and unknown keys are ignored
#

[Devices]
Count=1          # number of devices to configure
Dev_1=sdl1a0     # the name of a device
#Dev_2=sdl1a1    # the name of a device

# Specified here, these are applied to all devices and can be overridden for
# each individual board.
#
Access=CPE
Clock=Internal
KBaud=64
Flags=TX
#
# MTU=1500        # Maximum transmit IFrame length, default is 4096
# T391=10         # T391 value      5 - 30, default is 10
# T392=15         # T392 value      5 - 30, default is 15
# N391=6          # N391 value      1 - 255, default is 6
# N392=3          # N392 value      1 - 10, default is 3
# N393=4          # N393 value      1 - 10, default is 4

# Specified here, these set the defaults for all boards
# CIRfwd=16       # CIR forward    1 - 64
# Bc_fwd=16       # Bc forward     1 - 512
# Be_fwd=0        # Be forward     0 - 511
# CIRbak=16       # CIR backward   1 - 64
# Bc_bak=16       # Bc backward    1 - 512
# Be_bak=0        # Be backward    0 - 511

#
#
# Device specific configuration
#
#
#
# The first device is a Sangoma S502E
#
[sdl1a0]
Type=Sangoma     # Type of the device to configure, currently only
                 # SANGOMA is recognized
#
# These keys are specific to the 'Sangoma' type
#
# The type of Sangoma board - S502A, S502E, S508
Board=S502E
#
# The name of the test firmware for the Sangoma board
# Testware=/usr/src/frad-0.10/bin/sdl1a_tst.502
#
# The name of the FR firmware
# Firmware=/usr/src/frad-0.10/bin/frm_rel.502
#
Port=360         # Port for this particular card
Mem=C8           # Address of memory window, A0-EE, depending on card
IRQ=5            # IRQ number, do not supply for S502A
DLCIs=1          # Number of DLCI's attached to this device
DLCI_1=16        # DLCI #1's number, 16 - 991
# DLCI_2=17
# DLCI_3=18
# DLCI_4=19
# DLCI_5=20
#
# Specified here, these apply to this device only,
# and override defaults from above
#
# Access=CPE      # CPE or NODE, default is CPE
# Flags=TXIgnore,RXIgnore,BufferFrames,DropAborted,Stats,MCI,AutoDLCI

```

```

# Clock=Internal          # External or Internal, default is Internal
# Baud=128                # Specified baud rate of attached CSU/DSU
# MTU=2048                # Maximum transmit IFrame length, default is 4096
# T391=10                 # T391 value      5 - 30, default is 10
# T392=15                 # T392 value      5 - 30, default is 15
# N391=6                  # N391 value      1 - 255, default is 6
# N392=3                  # N392 value      1 - 10, default is 3
# N393=4                  # N393 value      1 - 10, default is 4

#
# The second device is some other card
#
# [sdla1]
# Type=FancyCard          # Type of the device to configure.
# Board=                  # Type of Sangoma board
# Key=Value               # values specific to this type of device

#
# DLCI Default configuration parameters
# These may be overridden in the DLCI specific configurations
#
CIRfwd=64                 # CIR forward      1 - 64
# Bc_fwd=16               # Bc forward      1 - 512
# Be_fwd=0                # Be forward      0 - 511
# CIRbak=16               # CIR backward    1 - 64
# Bc_bak=16               # Bc backward     1 - 512
# Be_bak=0                # Be backward     0 - 511

#
# DLCI Configuration
# These are all optional. The naming convention is
# [DLCI_D<devicenum>_<DLCI_Num>]
#
[DLCI_D1_16]
# IP=
# Net=
# Mask=
# Flags defined by Sangoma: TXIgnore,RXIgnore,BufferFrames
# DLCIFlags=TXIgnore,RXIgnore,BufferFrames
# CIRfwd=64
# Bc_fwd=512
# Be_fwd=0
# CIRbak=64
# Bc_bak=512
# Be_bak=0

[DLCI_D2_16]
# IP=
# Net=
# Mask=
# Flags defined by Sangoma: TXIgnore,RXIgnore,BufferFrames
# DLCIFlags=TXIgnore,RXIgnore,BufferFrames
# CIRfwd=16
# Bc_fwd=16
# Be_fwd=0
# CIRbak=16
# Bc_bak=16
# Be_bak=0

```

When you have built your `/etc/frad/router.conf` file, the only step remaining is to configure the actual devices themselves. This is only a little trickier than a normal network device configuration, you need to remember to bring up the FRAD device before the DLCI encapsulation devices. These commands are best hosted in a shell script, due to their number:

```

#!/bin/sh
# Configure the frad hardware and the DLCI parameters
/sbin/fradcfg /etc/frad/router.conf || exit 1
/sbin/dlcicfg file /etc/frad/router.conf
#
# Bring up the FRAD device
ifconfig sdla0 up

```

```
#
# Configure the DLCI encapsulation interfaces and routing
ifconfig dlci00 192.168.10.1 pointopoint 192.168.10.2 up
route add -net 192.168.10.0 netmask 255.255.255.0 dlci00
#
ifconfig dlci01 192.168.11.1 pointopoint 192.168.11.2 up
route add -net 192.168.11.0 netmask 255.255.255.0 dlci00
#
route add default dev dlci00
#
```

12.8.8. IPX (AF_IPX)

The IPX protocol is most commonly utilized in Novell NetWare^(tm) local area network environments. Linux includes support for this protocol and may be configured to act as a network end point, or as a router for IPX.

Kernel Compile Options:

```
Networking options --->
  [*] The IPX protocol
  [ ] Full internal IPX network
```

The IPX protocol and the NCPFS are covered in greater depth in the IPX-HOWTO (<http://linuxdoc.org/HOWTO/IPX-HOWTO.html>).

12.8.9. NetRom (AF_NETROM)

NetRom device names are nr0, nr1, etc.

Kernel Compile Options:

```
Networking options --->
  [*] Amateur Radio AX.25 Level 2
  [*] Amateur Radio NET/ROM
```

The AX25, Netrom and Rose protocols are covered by the AX25-HOWTO (<http://linuxdoc.org/HOWTO/AX25-HOWTO.html>). These protocols are used by Amateur Radio Operators world wide in packet radio experimentation.

Most of the work for implementation of these protocols has been done by Jonathon Naylor, jsn@cs.nott.ac.uk.

12.8.10. Rose Protocol (AF_ROSE)

Rose device names are rs0, rs1, etc. in 2.1.* kernels. Rose is available since the 2.1.* kernels.

Kernel Compile Options:

```
Networking options --->
  [*] Amateur Radio AX.25 Level 2
  <*> Amateur Radio X.25 PLP (Rose)
```

The AX25, Netrom and Rose protocols are covered by the AX25-HOWTO (<http://linuxdoc.org/HOWTO/AX25-HOWTO.html>). These protocols are used by Amateur Radio Operators world wide in packet radio experimentation.

Most of the work for implementation of these protocols has been done by Jonathon Naylor, jsn@cs.nott.ac.uk.

12.8.11. Samba - NetBEUI, NetBios, CIFS Support

Samba is an implementation of the Session Management Block protocol (SMB). It allows *Windows* and other systems to mount and use your disks and printers.

Samba and its configuration are covered in detail in the SMB-HOWTO (<http://linuxdoc.org/HOWTO/SMB-HOWTO.html>).

12.8.12. STRIP Support (Starmode Radio IP)

STRIP device names are `st0`, `st1`, etc.

Kernel Compile Options:

```
Network device support --->
[*] Network device support
....
[*] Radio network interfaces
< > STRIP (Metricom starmode radio IP)
```

STRIP is a protocol designed specifically for a range of Metricom radio modems for a research project being conducted by Stanford University called the MosquitoNet Project (<http://mosquitonet.Stanford.edu>). There is a lot of interesting reading here, even if you aren't directly interested in the project.

The Metricom radios connect to a serial port, employ Spread Spectrum technology and are typically capable of about 100kbps. Information on the Metricom radios is available from the Metricom Web Server (<http://www.metricom.com/>) web site.



Metricom went bankrupt, but another company might buy the technology and re-activate the web site.

At present, the standard network tools and utilities do not support the STRIP driver, so you will have to download some customized tools from the MosquitoNet web server. Details on what software you need is available on the MosquitoNet Software Page (<http://mosquitonet.Stanford.edu/software.html>).

A summary of configuration is that you use a modified **slattach** program to set the line discipline of a serial tty device to STRIP and then configure the resulting `st[0-9]` device, as you would for Ethernet with one important exception: for technical reasons, STRIP does not support the ARP protocol, so you must manually configure the ARP entries for each of the hosts on your subnet. This should not prove too onerous.

12.8.13. Token Ring

Token ring device names are `tr0`, `tr1` etc. Token Ring is an IBM standard LAN protocol that avoids collisions by providing a mechanism that gives only one station on the LAN the right to transmit at a time. A "token" is held by one station at a time and the station holding the token is the only station allowed to transmit. When it has transmitted its data, it passes the token onto the next station. The token loops amongst all active stations, hence the name "Token Ring".

Kernel Compile Options:

```
Network device support --->
[*] Network device support
....
[*] Token Ring driver support
< > IBM Tropic chipset based adaptor support
```

The configuration of Token Ring is identical to that of Ethernet, with the exception of the network device name to configure.

12.8.14. X.25

X.25 is a circuit-based packet switching protocol defined by the ITU-T (the Telecommunications Standardization Section of the International Telecommunications Union, a standards body recognized by Telecommunications companies in most parts of the world). An implementation of X.25 and LAPB are being worked on and recent kernels (from 2.1.*) include the work in progress.

Jonathon Naylor <jsn@cs.nott.ac.uk> is leading the development and a mailing list has been established to discuss Linux X.25-related matters. To subscribe, send a message to majordomo@vger.rutgers.edu with the text `subscribe linux-x25` in the body of the message.

12.8.15. WaveLan Card

WaveLan device names are `eth0`, `eth1`, etc.

Kernel Compile Options:

```
Network device support --->
[*] Network device support
....
[*] Radio network interfaces
....
<*> WaveLAN support
```

The WaveLAN card is a Spread Spectrum wireless LAN card. The card looks very much like an Ethernet card in practice, and is configured in much the same way.

You can get information on the WaveLan card on the ORiNOCCO (<http://www.orinocowireless.com/>) web site.

12.9. Cables and Cabling

Those of you handy with a soldering iron may want to build your own cables to interconnect two Linux machines. The following cabling diagrams should assist you in this.

12.9.1. Serial NULL Modem cable

Not all NULL modem cables are alike. Many null modem cables do little more than trick your computer into thinking all the appropriate signals are present and swap transmit and receive data. This is ok, but it means you must use software flow control (XON/XOFF) which is less efficient than hardware flow control. The following cable provides the best possible signalling between machines and allows you to use hardware (RTS/CTS) flow control.




Figure 12-2. The NULL-Modem Cabling

12.9.2. Parallel Port Cable (PLIP Cable)

If you intend to use the PLIP protocol between two machines, then this cable will work for you notwithstanding what sort of parallel ports you have installed.

Pin Name	pin	pin
STROBE	1*	
D0->ERROR	2	15
D1->SLCT	3	13
D2->PAPOUT	4	12
D3->ACK	5	10
D4->BUSY	6	11
D5	7*	
D6	8*	
D7	9*	
ACK->D3	10	5
BUSY->D4	11	6
PAPOUT->D2	12	4
SLCT->D1	13	3
FEED	14*	
ERROR->D0	15	2
INIT	16*	
SLCTIN	17*	
GROUND	25	25



- Do not connect the pins marked with an asterisk (“*");
- extra grounds are 18,19,20,21,22,23 and 24;
- if the cable you are using has a metallic shield, it should be connected to the metallic DB-25 shell at **one end only**.



A badly wired PLIP cable can destroy your controller card. Be very careful and double-check every connection to ensure you don't cause yourself any unnecessary work or heartache.

While you may be able to run PLIP cables for long distances, you should avoid it if you can. The specifications for the cable allow for a cable length of about 1 meter or so. Please be very careful when running long PLIP cables as sources of strong electromagnetic fields such as lightning, power lines and radio transmitters can interfere with and sometimes even damage your controller. If you really want to connect two of your computers over a large distance, you really should be looking at obtaining a pair of thin-net Ethernet cards and running some coaxial cable.

12.9.3. 10base2 (Thin Coax) Ethernet Cabling

10base2 is an Ethernet cabling standard that specifies the use of 50 ohm coaxial cable with a diameter of about 5 millimeters. There are a couple of important rules to remember when interconnecting machines with 10base2 cabling. The first is that you must use terminators at **both ends** of the cabling. A terminator is a 50 ohm resistor that helps to ensure that the signal is absorbed and not reflected when it reaches the end of the cable. Without a terminator at each end of the cabling, you may find that the Ethernet is unreliable or doesn't work at all. Normally, you should use “T pieces” to interconnect the machines, so that you end up with something that looks like:




Figure 12-3. 10base2 Ethernet Cabling

where the | at either end represents a terminator, the ===== represents a length of coaxial cable with BNC plugs at either end, and the T represents a “T piece” connector. You should keep the length of cable between the “T piece” and the actual Ethernet card in the PC as short as possible, ideally the “T piece” will be plugged directly into the Ethernet card.

12.9.4. Twisted-Pair Ethernet Cable

If you only have two twisted pair Ethernet cards and you wish to connect them, you do not require a hub. You can cable the two cards directly together. A diagram showing how to do this is included in the Ethernet-HOWTO (<http://linuxdoc.org/HOWTO/Ethernet-HOWTO.html>).

Appendix A. Where to Find Additional Documentation

Although we described the *MandrakeSecurity* interface, as well as reveal some details about network and security, this manual is simply not enough if you have no knowledge about some specific features we proposed. This is why we suggest links to various documents over the Web which will help you understand the actual mechanisms beneath *MandrakeSecurity*.

Bibliography

Documentation Related to the Back-end

Shorewall Documentation (http://www.shorewall.net/shorewall_quickstart_guide.htm).

The *Shorewall* application, on which *MandrakeSecurity* is based, is extensively documented. Look here for additional examples, and complete information on *MandrakeSecurity*'s background processes.

DNS Caching: ISC Bind (<http://www.isc.org/products/BIND/>).

DHCP (<http://www.isc.org/products/DHCP/>).

Proxy Web (<http://www.squid-cache.org/>).

SquidGuard URL Filtering (<http://www.squidguard.org/>).

DansGuardian URL Filtering (<http://dansguardian.org/>).

SNORT Intrusion Detection System (<http://www.snort.org/>).

Prelude Intrusion Detection System (<http://www.prelude-ids.org/>).

Miscellaneous Information about the Techniques Used in MandrakeSecurity

IP Network Address Translation (<http://www.suse.de/~mha/linux-ip-nat/diplom/>).

Linux FreeS/WAN: implementation of IPSEC & IKE for Linux (<http://freeswan.org>).

Windows 2000 / Windows XP - Freeswan VPN (<http://vpn.ebootis.de/>).

Appendix B. The GNU General Public License

The following text is the GPL license that applies to most programs found in **Mandrake Linux** distributions.

Version 2, June 1991 Copyright (C) 1989, 1991 Free Software Foundation, Inc. 59 Temple Place - Suite 330, Boston, MA 02111-1307, USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

B.1. Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software – to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps:

1. copyright the software, and
2. offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

B.2. Terms and conditions for copying, distribution and modification

- 0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only

if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

- 1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

- 2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:
 1. You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.
 2. You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.
 3. If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

- 3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:
 1. Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
 2. Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
 3. Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable.

However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

- 4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.
- 5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.
- 6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.
- 7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

- 8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.
- 9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

- 10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

- 11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.
- 12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

Appendix C. GNU Free Documentation License

C.1. GNU Free Documentation License

Version 1.1, March 2000

Copyright (C) 2000 Free Software Foundation, Inc. 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other written document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you".

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (For example, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, whose contents can be viewed and edited directly and straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup has been designed to thwart or discourage subsequent modification by readers is not Transparent. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML designed for human modification. Opaque formats include PostScript, PDF, proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or proces-

sing tools are not generally available, and the machine-generated HTML produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies of the Document numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a publicly-accessible computer-network location containing a complete Transparent copy of the Document, free of added material, which the general network-using public has access to download anonymously at no charge using public-standard network protocols. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has less than five).
- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.

- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section entitled "History", and its title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- K. In any section entitled "Acknowledgements" or "Dedications", preserve the section's title, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section entitled "Endorsements". Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section as "Endorsements" or to conflict in title with any Invariant Section.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections entitled "History" in the various original documents, forming one section entitled "History"; likewise combine any sections entitled "Acknowledgements", and any sections entitled "Dedications". You must delete all sections entitled "Endorsements."

6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, does not as a whole count as a Modified Version of the Document, provided no compilation copyright is claimed for the compilation. Such a compilation is called an "aggregate", and this License does not apply to the other self-contained works thus compiled with the Document, on account of their being thus compiled, if they are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one quarter of the entire aggregate, the Document's Cover Texts may be placed on covers that surround only the Document within the aggregate. Otherwise they must appear on covers around the whole aggregate.

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License provided that you also include the original English version of this License. In case of a disagreement between the translation and the original English version of this License, the original English version will prevail.

9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See Copyleft (<http://www.gnu.org/copyleft/>).

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

C.2. How to use this License for your documents

To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notices just after the title page:

Copyright (c) YEAR YOUR NAME. Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.1 or any later version published by the Free Software Foundation; with the Invariant Sections being LIST THEIR TITLES, with the Front-Cover Texts being LIST, and with the Back-Cover Texts being LIST. A copy of the license is included in the section entitled "GNU Free Documentation License".

If you have no Invariant Sections, write "with no Invariant Sections" instead of saying which ones are invariant. If you have no Front-Cover Texts, write "no Front-Cover Texts" instead of "Front-Cover Texts being LIST"; likewise for Back-Cover Texts.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.

